Real Robot Hand Grasping using Simulation-Based Optimisation of Portable Strategies

Frank Röthling

Der Technischen Fakultät der Universität Bielefeld
vorgelegt zur Erlangung des akademischen Grades
Doktor der Ingenieurswissenschaften

Mai 2007

Acknowledgement

This work was done in a perfectly organised environment, the Neuroinformatics Group, at the Faculty of Technology, Bielefeld University. To work out the complex subjects of this thesis would not have been possible without the conglomeration of state-of-the-art hardware and software that was available, respectively, was obtained when I was a member of this group. The good working atmosphere made me enjoying my time, and the fact that every colleague was addressable and cooperative almost at any time, helped to solve problems quickly.

First of all, I must thank the head of this group, Helge Ritter, for his support and his inspiring ideas. He has always set a good example for me in terms of his ways of seeing things positively and being motivated even when annoying work has to be done.

A special thanks goes to Jochen Steil from whom I learned a lot, not least because he always encouraged me to cope with tasks I had never done before. He is a professional, personable, and above all extremely competent advisor. He helped me improving my skills in writing, in presenting, and in many other matters.

Working with robot hardware can be very laborious considering maintenance. For their help in solving technical problems, I thank especially Oliver Lieske and Risto Kõiva. Finding ways for getting things going again was great fun with their assistance.

From Elena Carbone, I learned a lot about conducting experiments with people and statistical analysis. Thanks to you and to Stefan Krüger who assisted me greatly in performing the experiments and who was always willing to help.

Besides all other members of the Neuroinformatics Group, my thanks goes to Robert Haschke, Markus Henschel, and Jan Steffen. With their contributions and the fruitful collaboration with them, I was able to realise my ideas for the simulation-based optimisation strategy.

I also owe a debt of gratitude to Rich Walker for advancing my skills in the English language. In his function as technical director of the Shadow Robot Company, it was a great pleasure to develop with him different aspects of the five-fingered robot hand.

On a personal note, I would like to thank my parents for always supporting my decisions and Katrin for her love, encouragement, and assistance.

Contents

1 Introduction

1.1 Motivation

Providing robot hands with grasping capabilities is one of the great challenges of robotics. Different application areas can benefit from progress in this field of research. Nowadays, a large number of robots are employed in industry to perform grasping tasks by utilising simple two-fingered or three-fingered grippers. Most of them are integrated in an automated production line and repeat a sequence of motions, while performing only one kind of movement for opening and closing the gripper.

These kinds of manipulators accomplish their tasks reliably in highly structured assembly environments. Non-manufacturing robots employed in unstructured environments can rarely be found. The reason is that the potential fields of applications like households, aerospace, the medical area, or hazardous environments require robot hands that are capable of performing a much wider range of tasks.

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. But on the algorithmic side, robust and stable grasping of everyday objects is still a major challenge even for the best artifical robot hands available. To provide these more-or-less anthropomorphic robot hands with algorithms that realise such grasping capabilities, a promising approach is to mimic human grasp strategies.

The only grasp strategy of a newborn baby is a simple, gripper-like closing movement of all fingers as a reflex action. When the baby grows up, it learns to differentiate amongst single finger postures. An adult is able to perform a vast variety of postures with their dextrous hand. But the grasp postures people use in grasping everyday objects can be classified into a taxonomy consisting of only a few grasp types. The postures belonging to a grasp type people unconsciously choose from this taxonomy to grasp a particular object are optimised for the task and the object in a lifelong learning process.

1.2 Approach

Motivated by the strategies people utilise for grasping and for grasp optimisation, we propose strategies that are applicable to nearly all artificial hands and that are portable over different robot hand setups. These strategies provide robot hands with capabilities for grasping a variety of everyday objects in unstructured environments.

For this purpose, we use an approach that mimics human grasping and comprises (i) the identification of a grasp model consisting of different grasp components and grasp phases; (ii) the differentiation of significant hand postures and grasp types; and (iii) the determination of a strategy, based upon the identified grasp model, for applying a grasp to a target object.

To provide portability, the grasp strategy and the grasp types have to be defined taking into consideration the capabilities of existing robot hands. By implementing the grasp strategy and realising

the grasp types on two completely different robot hand setups, including artificial hands that differ in most characteristics, we prove the portability of our approaches.

For evaluating the suitability of different grasp types, the success of grasp optimisation strategies, and the grasp capabilities of different robot hands, we propose a benchmark system consisting of a variety of everyday objects and an assessment test. Because these objects are different in their properties, a few grasp types realised cannot provide optimal solutions for all objects. Similar to human grasping, the most suitable grasp type for a particular object can be chosen, and its postures have parameters that can be usefully optimised.

To this end, we propose an optimisation strategy capable of optimising different characteristics of object-specific grasp postures and which is supported by an experiment on human grasping. For optimising the grasp postures of both real artificial hands, we use a simulation environment containing an exact mapping of the robot hands, the benchmark objects, and the grasps. There are two reasons for employing a simulator for grasp optimisation. Firstly, most real robot hand setups do not possess suitable sensors for evaluating the grasp quality needed to optimise the grasp postures. And secondly, one of the two optimisation steps we propose needs around 1000 grasp trials to optimise one grasp posture. It is not reasonable to perform this with a real robot hand.

By analysing a number of extensive grasp experiments utilising our benchmark system, we demonstrate the suitability of our grasp strategy, the realisability of the proposed grasp types on the different robot hand setups, and the portability together with the improvement capabilities of our simulation-based optimisation strategy.

1.3 Contributions

The contributions of this thesis in the areas of human and robot grasping are subdivided according to the following chapters:[1]

Chapter 2: Based on the most significant existing studies on human grasping, we develop a *model of human grasping*. By partly adopting this model, a *model of robot grasping* is proposed, and the correlations within grasp components and grasp phases of both models are highlighted. Our approach to a *robot grasp strategy* is based on these models and differs from most other biologically motived approaches mainly in that, besides the pre-grasp posture, it stresses the *target grasp posture*. In a *taxonomy for robot grasping* we define *four grasp types* that can be realised by most anthropomorphic robot hands possessing at least three fingers.

Chapter 3: The grasp strategy is *implemented* on two different robot hand setups. By *developing* a palm and reconfiguring the three fingers of a gripper, we obtained an anthropomorphic three-fingered hand which is the core piece of one of the setups. The second setup was utilisable for robot grasping after purchase and *implementation* of a dextrous anthropomorphic five-fingered hand. The developmental process for realising the grasp types[2] on both setups is described, and *general development rules* are proposed.

Chapter 4: Grasp types that are realised on any setup including a robot hand mounted on a robot arm can be evaluated with our *benchmark system*. This system consists of 21 *benchmark objects* and a *benchmark test* determining rules for conducting an experiment that results in

[1] For definitions of terms used in this thesis, the reader is referred to the following chapters and to the glossary at the end of the thesis.

[2] The four grasp types defined by our taxonomy plus one additional.

a grasp success rate for each investigated combination of grasp, grasp strategy, and robot hand setup. The *evaluation* of the realised grasp types in a preliminary experiment leads to the determination of the most suitable grasp type for each benchmark object.

Chapter 5: To substantiate ideas for optimising object-specific grasps gained during the performance of the preliminary experiment, we conducted an *experiment on human grasping*. In this, we investigated the *contact strategy* and the *contact simultaneity* in human grasping. For this purpose, we define four different *measures of contact simultaneity* and propose a more reliable measure for *hand opening* than the commonly used "grasp aperture". The experiment leads to the *result* that humans strive for contact simultaneity when they grasp an object.

Chapter 6: Supported by this result, we propose a *grasp optimisation strategy* that, in its first step, *optimises the pre-grasp posture* of a grasp for contact simultaneity. The second step of the optimisation strategy *optimises the target grasp posture* to give the best closure trajectory of the thumb by applying an *evolutionary algorithm*. This optimisation is performed with a physics-based grasp simulator after implementing both robot hands, all benchmark objects and all standard grasps. For determining the quality of a simulated grasp and for providing a fitness value to the evolutionary algorithm, we propose a *grasp stability measure*.

Chapter 7: To evaluate the optimisation strategy we performed several *grasp experiments* based on our benchmark system with both of the real robot hands. The object-specific grasps were *benchmarked* against each other before optimisation, after the first optimisation step, and after the second optimisation step. A detailed *analysis* of the simulated and real grasps are presented and the optimisation strategy is evaluated. One additional grasp experiment leads to a *comparison* between both robot hands in terms of grasping capabilities.

Chapter 8: The accomplishments are summarised, conclusions about the optimisation strategy are drawn, and an outlook on potential improvements is presented.

2 From Human to Robot Grasping

A variety of contributions in the research area of robot grasping has established a theoretical framework for grasp analysis, synthesis, and simulation. Bicchi and Kumar [2000] present a survey of analytical approaches[1] which have dominated the field for a long time. More recently, biologically motivated approaches comprising knowledge-based approaches, control-based approaches, and behaviour-based approaches have been considered to a greater extent. But still there is a gap between the theoretical promise and the practical delivery [Okamura et al., 2000]. To close this gap, we use a biologically motivated approach that provides the opportunity to easily synthesise grasps for successful grasping of everyday objects on different robot hand setups.

The following Section 2.1 gives an overview of biologically motivated approaches which consider the capabilities of the grasping system, and analytical approaches which, in contrast, rather focus on contact points or contact regions of the target object. The models of human and robot grasping that provide the basis of our robot grasp strategy are introduced in Section 2.2. The grasp strategy defines a sequence of steps determining the way a grasp is applied with a robot hand. After reviewing the most significant grasp classifications in Section 2.3, we propose different grasp types that are realisable by nearly all artificial hands and are classified into a taxonomy for robot grasping.

2.1 Grasp Synthesis

The study of grasp synthesis can be categorised into two broad groups. In analytical approaches, a *grasp* is formally defined as a set of contact points on the surface of the target object together with friction cone conditions, whereas being independent of the robot hand under investigation [Murray et al., 1994]. Using this grasp definition, grasp synthesis is the problem of deciding where to place the grasp contacts on the surface of an object. For grasp execution, a grasping system has to be able to precisely reach these contact points. In simulation and highly structured environments this kind of grasp definition can be useful, but when grasping real world objects with a robot hand, uncertainties, especially in the object's location and orientation, restrain the manipulator from reaching these contact points.

To realise more flexibility and robustness in grasping, several authors have proposed organising robot grasping in a more holistic fashion loosely motivated by the way humans grasp. Such biologically motivated approaches are primarily based on empirical studies of human grasping and manipulation. In human grasping, planning the precise location for finger placement is not needed, as humans can anticipate object behaviour during the interaction between hand and object [MacKenzie and Iberall, 1994]. Humans select an object-specific pre-grasp posture as one of a few prehensile hand postures [Cutkosky and Howe, 1990], and the grasp itself is carried out by comprehensively closing the fingers and evaluating the tactile feedback. Though differing in detail, *grasp* definitions in the group of biologically motivated approaches imply different grasp phases, a pre-grasp posture, and a grasp closing strategy.

The following two sections review promising approaches of both groups.

[1] Sometimes analytical approaches are referred to as theoretical approaches or model-based approaches.

2.1.1 Analytical Approaches

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 [Borst et al., 2002]. Since for most objects the optimal set of contact points is not realisable for most robot hands [Borst et al., 2003], a compromise between planning optimal grasps and the constraints of the hand kinematics has to be found.

Borst et al. [1999] use a heuristic generation of grasp candidates (sets of contact points on the target object) for grasping with the four-fingered DLR Hand [Butterfass et al., 2004]. After choosing the first contact point, the other three points are determined by using a geometric algorithm including some kind of arbitrariness. After contact points are moved away from the edges to achieve a more robust grasp, the best grasp candidate is chosen by utilising a grasp quality measure. This measure is based on the definition of force-closure and is determined by the magnitude of the largest worst-case disturbance wrench that can be resisted by a grasp of unit strength (for more explanation, see Section 6.2.1). Borst et al. [2005] give an overview of different grasp qualification methods and discuss the quality measure used in more detail.

From the found set of contact points, a feasible hand posture has to be determined in a subsequent calculation step [Borst et al., 2004]. The authors use a generic method in which an objective function with penalty terms has to be minimised. Depending on the initial hand configuration, this minimisation method converges into different local minima, and a solution space with many equal solutions results. As Borst et al. sum up, how to optimise in this solution space is an open question.

Two major problems concerning real hand kinematics led to diverse approaches of determining ”independent regions of contact” for each finger. These problems are (i) the optimal set of contact points is often not realisable; and (ii) the accuracy limitations of robot hands are responsible for uncertainty in finger positioning. One approach, leading to robustness to contact positioning errors, maximises the size of the independent contact regions, while force-closure is maintained as long as one contact is placed in each region. Early work in this area was done by Nguyen [1988]. The author outlined an algorithm for directly constructing a complete set of all force-closure grasps based on the independent regions of contact. Ponce et al. [1993] extended this approach to the synthesis of three- and four-finger force-closure grasps of polyhedral objects.

To accommodate constraints and errors in contact placement, Pollard [1996] determines regions around contacts of a predefined ”good” example grasp. These regions are similar to the robust regions found by Nguyen, and can be used to derive grasps for a variety of target object geometries. Each contact displaced anywhere in its corresponding region still leads to a force-closure grasp. In Pollard [2004] this approach is extended for considering multiple contacts in each independent region. Although the kinematics of the mechanical system were ignored, Pollard argues that grasps based upon these contact regions are likely to be achievable, because the kinematic configurations of the mechanism are similar to that of the example grasp.

In summary, analytical approaches to grasp synthesis find optimal contact points or contact regions. To assure that a kinematic configuration of the mechanism to reach these points can be found, remains a problem. Promising approaches additionally use some kind of knowledge about feasible grasps.

2.1.2 Biologically Motivated Approaches

Biologically motivated approaches use knowledge about feasible and useful grasp prototypes and parameters like orientation and position of the grasping hand. This knowledge can be encoded in primitive controllers, and complex grasping behaviour is realised by combining different controllers. The effectiveness of the resulting grasp process has to be verified by its execution on the real grasping mechanism or in simulation.

Miller et al. [2003] used a modelled grasp scene in the simulation environment GraspIt [Miller and Allen, 2004] for planning and evaluation of robot grasping. Objects are modelled as a set of shape primitives. By defining rules to determine the position, the orientation and the pre-grasp posture of the hand for each shape primitive, a set of different grasp preconditions is generated. The best grasp is evaluated by using a quality measure after executing the grasp in simulation. As denoted in Miller et al. [2005], the drawback of this approach is that, since it only considers a subset of the possible grasps, it may miss a better possibility. Pelossof et al. [2004] present an approach that considers a broader range of possible grasps by using support vector machine (SVM) regression to predict grasps for known and unknown objects for grasping with a three-fingered Barrett Hand [Townsend, 2000]. To reduce the search space, only two from ten possible grasp parameters were randomly chosen. The training set consisted of 1,600 grasps (100 random combinations of the two parameters for each of 16 grasp starting positions) for grasping each of 9 different objects. Pelossof et al. show that reasonable grasps of novel objects can be predicted, and although the regression does not always favour the simulated best grasp, it typically chooses grasps that still perform well.

Control-based and behaviour-based approaches to robot grasping are further promising biologically motivated approaches. Complex behaviour is created by sequencing and combining primitive controllers or behaviours. Platt et al. [2005] distinguish two different reach and two different grasp controllers. One grasp controller realises a three finger grasp, while the other grasp controller combines two physical contacts into a ”virtual finger” [Platt et al., 2003]. 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 [Platt et al., 2006]. This autonomous exploration is executed on real robot hardware, while objects are grasped with a Barrett Hand.

We also use a biologically motivated approach, but the major difference to those introduced in this section is that in we lay more stress on the target grasp posture as examined in Section 2.2.

2.2 Grasp Strategies

Human grasping has been investigated in many studies. Most of them identify a model comprising two grasp control components. Both of these components can be divided into different phases. The grasp strategy determines the chronological order in which these phases are executed and how the grasp postures correlate to them.

The next sections develop a model of human grasping based on the most significant existent studies. It is compared to the model of robot grasping we propose, before our approach to a robot grasp strategy is introduced in Section 2.2.4.

2.2.1 Components of Human Grasping

The model of human grasping can be divided into two distinct sensorimotor control components: *transportation* and *manipulation* [Jeannerod, 1981]. The *transportation* component controls the

arm movements to direct the hand from an initial position towards the target object and further to a lifted position. The *manipulation* component controls the hand movements to pre-shape, grasp, and stabilise the object.[2] Human grasping is segmented into these two components which are independently controlled in parallel [Jeannerod, 1984].

Other authors, like Fan et al. [2005], identify a third component as an individual channel of prehension movements: the hand orientation. In his investigation of human grasping motor schemas, Arbib [1981] differentiates ballistic movements (transportation) from finger adjustment (manipulation) and hand rotation. He relates the latter two by combining them into one motor schema for grasping. In contrast, Desmurget et al. [1996] relate hand rotation to arm transportation and advise that they do not constitute independent channels. Another hypothesis of Gentilucci et al. [1996] suggests that hand orientation is a result of an independent process of the two control components (transportation and manipulation).

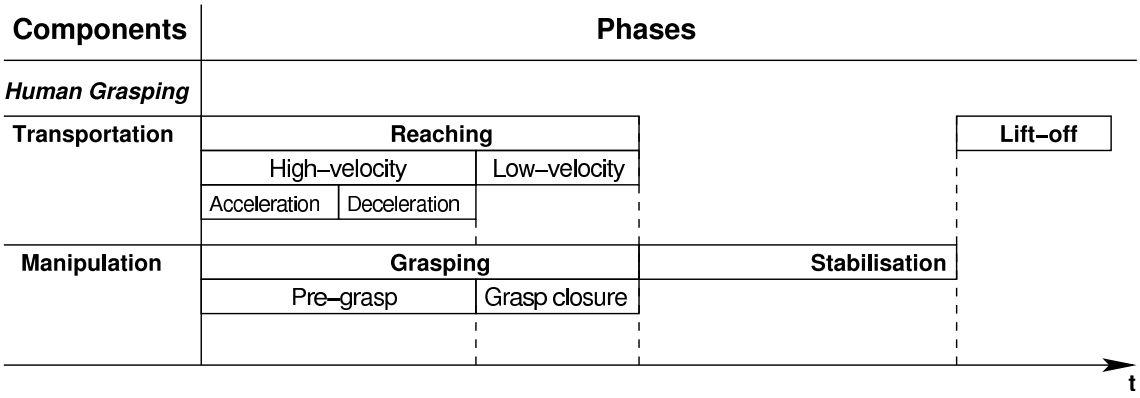

Figure 2.1: Chronological order of human grasp phases. Dashed lines indicate correlated points in time. Widths of the blocks have no quantitative meaning.

However, we support the commonly used division of human grasping into two components where the orientation belongs to the transportation component. This model is pictured in Figure 2.1 and can be correlated with robot grasping as Section 2.2.4 reveals.

2.2.2 Phases of Human Grasping

The two sensorimotor components can be divided into serial ordered phases. We propose the following distinction as depicted in Figure 2.1. The transportation component consist of a *reaching* and a *lift-off phase*. The manipulation component is also divided into two phases: *grasping* and *stabilisation*. The subdivision of the reaching and the grasping phase was introduced in the grasping literature by Jeannerod [1981]. In contrast to our approach, Jeannerod equated the terms transportation with reaching and manipulation with grasping, while disregarding other phases of human grasping.

As a person reaches for an object, the arm accelerates, reaches its maximum velocity, decelerates, and re-accelerates near to the target object [Jeannerod, 1984]. The point of re-acceleration divides the reaching phase into a *high-velocity* and a *low-velocity* phase. Parallel to the reaching phase of the transportation component, the hand pre-shapes into a pre-grasp posture suitable

[2] Manipulation in terms of moving the object between the fingers (also known as ”dynamic grasping” [Iberall and Lyons, 1984] or ”finger gaiting” [Huber and Grupen, 2002]) is outside the scope of this thesis.

for the interaction and then begins to close in anticipation of contact with the object. Thus, the grasping phase of the manipulation component is subdivided into the phases of *pre-grasp* and *grasp closure*. The point of the maximum hand opening (in most studies referred to as peak grasp aperture; for more explanation see Section 5.2) is defined as the beginning of the grasp closure phase. Jeannerod [1984] found that this point corresponds to the point of re-acceleration of the arm movement. Additionally, the low-velocity phase (transportation component) and the grasp closure phase (grasping component) end simultaneously. Therefore, these two phases are coordinated in time (indicated by dashed lines in Figure 2.1). In studies conducted by Jeannerod, this coordination was found in conditions where visual feedback from the moving limb was present or absent.

The last phase of the transportation component of human grasping is the *lift-off* phase. Usually the ends of both the reaching phase and the grasping phase were defined to be the beginning of the lift-off phase. Thus, in most human grasping experiments the end of the grasp closure phase was determined by the first movement of the target object (for example Zaal and Bootsma [1993]) or by using a finger velocity threshold (for example Smeets and Brenner [2001]). But these studies disregard the time needed by the fingers for force application after contacting the object and before lifting it. Biegstraaten et al. [2006] found that fingers spend about 200 ms in contact with the object before it starts to move. While Biegstraaten et al. suggest that this time period is needed for a gradual transition between the grasping phase and the lift-off phase, Weir et al. [1991] identify it as the ”finger-object interaction phase”. Weir et al. found that the duration of this phase increases for objects with higher weights and for all objects whose weights are unknown. They conclude that there are at least two motor control phases involved in the manipulation component of prehension, one for making contact with the object and the other for finger-object interaction.

For quantifying grasp phases and finger contact times, a precise definition of the end of the grasp closure phase is necessary (see Section 5.4). We propose to identify the end of grasp closure as the first point in time when all grasping fingers touch the target object. This point also determines the beginning of the *stabilisation* phase, which ends with the first object movement in the lift-off direction (right dashed line in Figure 2.1).

2.2.3 Postures of Human and Robot Grasping

When the human hand reaches out to grasp an object, it pre-shapes into a posture that closely resembles the posture taken on in contacting the object [Iberall and MacKenzie, 1990]. This *pre-grasp posture* depends on visually determined estimation of the object’s size and shape. After pre-shaping the hand, the grasp is applied in the subsequent closure phase. When all grasping fingers touch the target object, the hand adopts the *grip posture*.

Robot grasp approaches that try to mimic human grasping define rules to generate pre-grasp postures or choose a pre-grasp posture from a predefined set of prototypes. After the hand is pre-shaped, a simple closing mechanism is applied to realise the grip posture. In our approach, we choose between different grasp prototypes[3] each comprising a pre-grasp posture and a *target grasp posture*. The latter is the posture the hand adopts when no object is located between the fingers.

In the following, the three different postures are discussed in detail, before the importance of the *thumb posture* is pointed out.

[3] The basic grasp types are defined in Section 2.3.2.

Pre-Grasp Posture

The majority of studies that investigated *pre-grasp postures* analysed the ”aperture” of the hand as a simplification of the hand opening (see Section 5.2). The hand *aperture* is defined as the distance between the tips of the thumb and the index finger. When the hand reaches its maximal aperture, also called ”peak grasp aperture”, the pre-grasp phase ends and the pre-grasp posture is realised. Jeannerod [1981] found that the size of the peak grasp aperture is a function of the anticipated size of the object, i.e. it is larger when the movement is directed at a large object. In any case, it is always larger than the actual size of the target object would require.

Several investigations revealed that the size of the pre-grasp posture is independent of object position (transportation component) and of several kinds of perturbations. Paulignan et al. [1991b] showed that peak grasp aperture is not effected by object location. Even in perturbed trials, when an instantaneous change of the target location occurred after hand movement initiation in direction of the primary object position, the peak grasp aperture had the same size as under normal grasp conditions. In their experiment, there were only small differences in distances of object locations to the participants. But other studies involving larger distances show similar results (for example, Chieffi and Gentilucci [1993]; Gentilucci et al. [1991]. In another study, Paulignan et al. [1991a] found that there is no difference in peak grasp aperture whether the size of the target object is realised before the hand reaches out to grasp the object or afterwards. Even a perturbation in the form of switching the object size just before initiation of the grasp phase leads to the same peak grasp aperture as under normal conditions.

To summarise, these findings support the assumption that the size of the pre-grasp posture is not dependent on any terms influencing the transportation component of human grasping. Transferring these results to robot grasping means that in developing and optimising the pre-grasp posture, the only constraints to consider are intrinsic properties of the target object such as size and shape.

An experiment in which the peak grasp aperture was different when compared with normal grasp conditions was conducted by Wing et al. [Wing et al., 1986]. One requirement was to grasp the target object blind, and another one was to grasp it very fast. In both cases, the peak grasp aperture increased. This is caused by the increased inaccuracies in the relative location of the hand to the target object. In robot grasping, we also face these kinds of inaccuracies. Usually there are no requirements on velocity, but the whole grasp process cannot be visually guided (as further explained in Section 2.2.4). Thus, robot grasping is a kind of blind grasping. For the development of pre-grasp postures, that means that it is reasonable to achieve a hand opening as large as possible, while still considering a finger formation resembling the object’s shape. This increases the possibility that the grasping fingers surround the object which is a major requirement for successful grasping.

Grip Posture and Target Grasp Posture

When a human hand reaches its *grip posture*, all finger pads are mostly arranged into the same plane [Iberall and MacKenzie, 1990]. If the closing fingers would not be obstructed by the target object, they would approach the same target point. This assumption is supported by perturbation studies of rapid grasps with the thumb and the index finger conducted by Cole and Abbs [1987]. The authors observed that subjects consistently brought the finger pads into contact although that was not part of the task instructions.

These findings can lead to the assumption that closure movements of the grasping fingers are correlated and very similar. Such a behaviour is adopted in most biologically motivated approaches to robot grasp synthesis (for example, Lyons [1985]; Wren and Fisher [1995]; and those introduced

in Section 2.1.2) in which only a simple closing mechanism is applied to realise the grip posture. This mechanism synchronously closes the fingers around the object and is independent of the selected pre-grasp type.

But although humans seem to apply a simple grasp closure strategy [Cutkosky and Howe, 1990], there are great differences compared to robot grasping. Human hands are endowed with compliant finger pads and an enormous number of tactile sensors that are used to adapt the final grip posture in a very fast and sophisticated closed-loop control. Current robot hand systems are far from realising these capabilities. Because a sophisticated adaptation of finger closure trajectories at the time when touching the target object is not realisable in robot grasping yet, these finger trajectories have to be optimised beforehand.

For this purpose, we define a *target grasp posture* that is reached when the fingers are not obstructed by an object. In determining a different target grasp posture for each grasp type, different closing behaviours of each grasping finger can be realised.

Importance of the Thumb Posture

The grip posture functionally has to provide appropriate forces to pick up the target object, it has to supply necessary friction, and it has to establish stability. When grasping an object lying on a flat desktop, at least two opposing forces are necessary.[4] Forces exerted by the fingers, except for the thumb, in most cases have similar directions. The thumb receives an exceptional position in that it is responsible for providing the opposing force by itself.

To this end, the thumb of a human hand is endowed with associated musculature that allows larger forces compared to the other fingers [Flatt et al., 2000]. A very articulated trapeziometacarpal saddle joint basically facilitates the capacity of opposing the thumb to all four fingertips [Marzke, 1992]. By utilising this saddle joint, thumb motions such as flexion / extension and abduction / adduction are possible. Additionally, a rotational movement is coupled to flexion / extension due to ligaments [MacKenzie and Iberall, 1994]. This flexibility provides the human hand with its dexterity and allows a large number of different *thumb postures*.

2.2.4 Approach to a Robot Grasp Strategy

The grasp strategy we propose is based on our **model of robot grasping**. This model is strongly related to the model of human grasping depicted in Figure 2.1. The two models are compared in Figure 2.2.

We propose the same simplified distinction into two sensorimotor control components, as Jeannerod [1981] proposes for human grasping. These components also consist of two major phases each: *reaching* and *lift-off* in the case of the *transportation* component and *grasping* and *stabilisation* in the case of the *manipulation* component. The transportation component of a robot grasp system drives the joints of the robot arm to control the position and orientation of the hand. The manipulation component drives the joints of the hand to adopt a specified posture and to exert desired forces.

Reaching is subdivided into approach and placing. The *approach* phase corresponds to the *high-velocity* phase of human grasping. A coarse approach to the target object is realised in the *gross-motion* phase, while in the *fine-motion* phase the hand is positioned over the object and oriented along its main axes. This fine positioning can rely on local visual feedback provided by a wrist

[4] Balancing a tray or grasping a suitcase at its handle requires forces in only one direction.

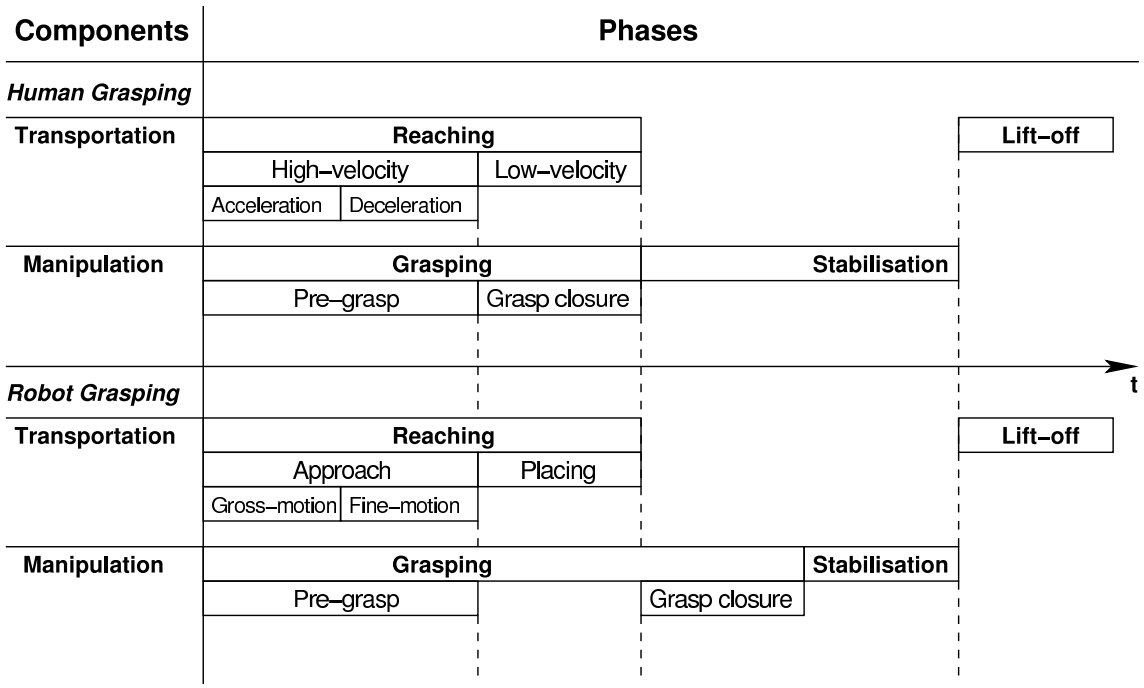

Figure 2.2: Chronological order of human and robot grasp phases. Dashed lines indicate correlated points in time. Widths of the blocks have no quantitative meaning.

camera, like that used in the TUM Hand setup described in Section 3.1.1. At the end of this phase, the hand reaches the *pre-grasp position*. With the hand in pre-grasp posture, the *placing* phase is entered which relates to the *low-velocity* phase of human grasping. In this phase, the hand is moved towards the *grasp position*.

A major difference between the models proposed for human and robot grasping is the time displacement of the *grasp closure* phase. In contrast to human grasping, where the fingers start to close while the hand approaches the target object, the grasp closure phase starts as soon as the reaching phase ends. In human grasping, the hand in pre-grasp posture is open wider than is necessary to enclose the target object. This is useful to compensate for inaccuracies in hand transportation. When the hand approaches the object, the actual required size of hand opening can be anticipated, and grasp closure can start before the object is reached. A robot system that incorporates visual surveillance to realise real-time control of arm and hand could mimic that human behaviour. But most existing robot systems (like ours) are not that sophisticated. Uncertainties about object position and orientation and inaccuracies in hand positioning and control can best be compensated by closing the grasp not before the grasp position is reached.

The Grasp Strategy

Based on this model of robot grasping, the grasp process can be divided into a sequence of five control phases while combining the transportation and manipulation components: (1) *Approach/Pre-grasp phase*; (2) *Placing phase* (3) *Grasp closure phase*; (4) *Stabilisation phase*; (5) *Lift-off phase*. As an advantage, a serial order of control phases simplifies the control needs, and robot grasping can be realised in a state-machine-like fashion [Steil et al., 2003].

Before the grasp process can be executed, the *grasp* g to be applied has to be chosen. We define four different *grasp types* t: *all finger precision* t_1, *two finger precision* t_2, *power* t_3, and *two*

finger pinch t_4. These grasp types constitute our grasp taxonomy as introduced in Section 2.3.2. The grasp type t defines the number of grasping fingers and qualitatively determines the target touch areas of the hand. For nearly every existing robot hand, having at least three fingers, a *standard grasp* of each grasp type t can be developed. Besides these (g_1 to g_4), we determine a fifth standard grasp, the *three finger special* g_5, as described in Section 3.2.3. Additionally, in our robot hand setups, we can choose between object-specific grasps that are optimised by the optimisation strategy proposed in Section 6.3.

Each *grasp* g comprises a *pre-grasp* and a *target grasp*. Pre-grasp and target grasp 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 parameterisations of the following characteristics are associated: a relative *position* $\boldsymbol{p}$ (3 DOF) and *orientation* $\boldsymbol{o}$ (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. The determination of these parameters is described in Section 3.2.1.

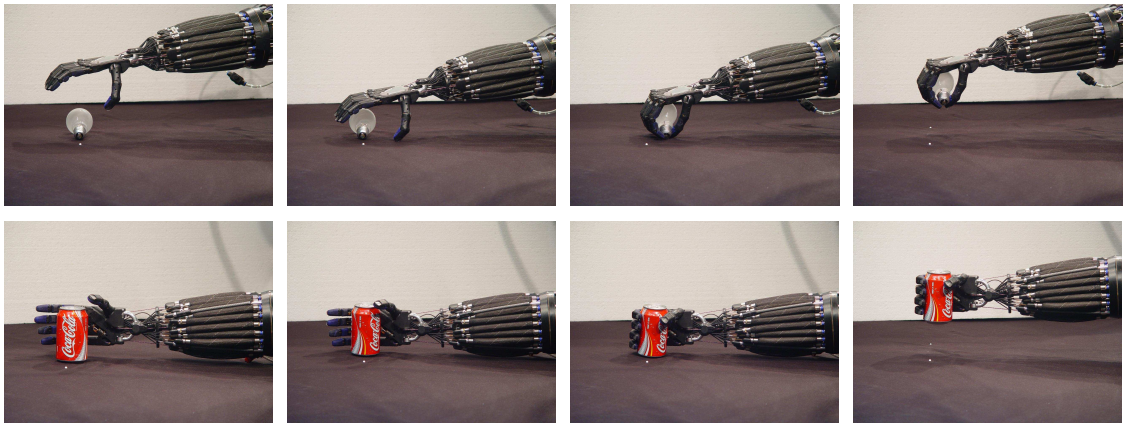

(a) Pre-grasp position. (after step 1)
(b) Grasp position. (after step 2)
(c) Object grasped. (after step 3 and 4)
(d) Object lifted. (after step 5)

Figure 2.3: The light bulb is grasped with the Shadow Hand by utilising our grasp strategy (upper row). This strategy can also be applied to grasping an object from other directions (lower row).

Based upon these definitions and our model of robot grasping, the **grasp strategy** comprises the following steps which are illustrated in Figure 2.3:

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. *Stabilisation phase*: Wait until fingers exert sufficient forces on the object.
5. *Lift-off phase*: Move hand to pre-grasp position.

We developed this grasp strategy for the purpose of grasping objects lying on a desktop from above, but our strategy can also be applied for other grasping tasks, just by determining an appropriate position $\boldsymbol{p}$ and orientation $\boldsymbol{o}$. One example is the task of grasping objects from the side as

pictured in Figure 2.3 (lower row). Here, the beverage can is grasped with the same grasp type, the power grasp t_3, as the light bulb (upper row). There is even no difference in the approach direction relative to a hand coordinate system.

2.3 Grasp Taxonomies

Robot hands which are built to mimic the dexterity of human grasping have to be endowed with a large number of degrees of freedom. A robot hand possessing motion capabilities similar to that of a human hand, like, for example the 20-DOF Shadow Hand (described in Section 3.1.2), provides an enormous set of potential hand configurations. To reduce the difficulty in determining the hand configurations suitable for robot grasping, a classification of grasp types is required.

The next section reviews various significant grasp taxonomies of different research fields. Most of them do not consider the applicability of their grasp types to robot grasping. Because these taxonomies are too detailed, too basic, or some grasp types that they cover cannot be realised with common robot hands, we propose a *taxonomy for robot grasping* in the final section of this chapter.

2.3.1 Existing Taxonomies

Classifying hand postures used for prehension for a long time has been an interest of fields such as hand surgery, rehabilitation, or the design of prosthetic devices. More recent application areas of grasp taxonomies are robotics and computer animation. Existing classifications lay different importances on the role of the hand, the object, and functional characteristics in prehension.

Object Shapes, Hand Surfaces, and Hand Shapes

Early approaches to categorising human grasps can be found in the area of rehabilitation and the design of prosthetic devices. Most of them use simple association of grasps to object or hand shapes, or to hand surfaces.

The purpose of a taxonomy developed by Schlesinger [1919] was to provide a basis for classifying prehensile functionality of prosthetic hands. At that time, substitute devices were needed in large quantity because of injuries from World War I. In a version summarised by Taylor and Schwarz [1955], the taxonomy comprises six basic types of prehension. These grasp types have symbolic names that relate to object shapes (cylindrical, spherical), hand surfaces (palmar, fingertip, lateral), and hand shapes (hook). Other classical taxonomies from the anthropological and medical fields use similar names and images. McBride [1942] distinguished grasps according to the hand surfaces used, while Griffiths [1943] suggested a taxonomy based on object shapes alone. The terms used in these taxonomies represent simple descriptions of the grasps. Functional characteristics, directions of applicable forces, or the purposes the grasps are appropriate for are not expressed.

Hand Potentials

In his study, Napier [1956] was interested in a fundamental analysis of the potentials of the hand as a whole. He was inspired by the need for a system of disability evaluation of the hand and divided the prehension into two basic patterns: precision grasp and power grasp. Depending on the purpose of the grasp, the same object can be held with either of the two grasps. The power grasp

provides large areas of contact between the grasped object and the surfaces of the fingers and the palm. Therefore, it is chosen when stability and security are necessary. If dexterity and sensitivity are more important, the precision grasp is used in which the object is held between the volar surfaces of the fingers and the thumb (usually the pads of the fingertips). These two grasp categories distinguish between major capabilities of the human hand. As a prehensile task can combine precision and power requirements, these two categories can be interpreted as characteristic concepts for prehension which are not mutually exclusive.

Lyons [1985] defined three grasps which he derived from studies of human reaching and grasping and from inspection of assembly domain tasks. Like Napier, he distinguished between a precision grasp and a power grasp which he called ”encompass grasp”. In his precision grasp, the object is held between the fingertips only, while a third grasp type, the lateral grasp, is characterised by all volar surfaces of the fingers for being contact areas.

Function of the Fingers

Arbib et al. [1985] classified the functions of the fingers respectively to the task they have to accomplish in grasping an object. These functions essentially are directions of the forces that the fingers exert on the target object. Grasping a mug, for example, requires a downward force from above the handle, an upward force from within the handle, and potentially a third force to stabilise the handle from below. The concept of a ”virtual finger” is used to categorise the different tasks the ”real” fingers have to perform, and fingers which have the same task are combined to one virtual finger. The number of the fingers that are combined might be different, even when the task requirements are identical. For example, the upward force from within the handle can be applied by one, two, three, or even four fingers (depending on the size of the handle). Instead of coping with many degrees of freedom, for instance in the case of a five-fingered robot hand, a task can be described by two or three virtual fingers. Combined fingers can be handled identically, which reduces the control needs.

Classification of Opposing Forces

Prehensile postures are constrained by the way the hand can apply opposing forces around an object for a given task [Iberall et al., 1986]. Therefore, Iberall [1987] proposed categorising these postures into three basic methods called oppositions: pad opposition, palm opposition, and side opposition. Pad opposition describes a hand posture where the thumb pad and the finger pads oppose each other. In palm opposition, forces exerted on the target object originate from the palm and the fingers. In side opposition, an object is held between the thumb pad and the side of the index finger, or it is fixed between the sides of two fingers. A grasp posture can consist of several of these oppositions. In Iberall [1997], 59 different grasp postures are divided into palm, pad, and side oppositions, and combinations thereof. The hand parts exerting opposing forces on the target object are mapped into virtual fingers according to their function as suggested by Arbib et al. [1985]. The relation between other classifications, like Schlesinger [1919], Napier [1956], Iberall and Lyons [1984], and Cutkosky and Wright [1986], and the opposition types combined with their mapping to the correlative virtual fingers is compared in Iberall [1987].

Task Requirements and Object Geometry

In the robotics literature, the best known taxonomy was presented by Cutkosky and Wright [1986]. It was a result of a study of the grasps used by machinists in a small batch manufacturing oper-

ation. Basing on the precision/power dichotomy of Napier, power grasping is subdivided into nine different grasp types and precision grasping into seven grasp types. To describe these types, Cutkosky and Wright use terms of many kinds of relations like object shapes (e.g. sphere), hand surfaces (e.g. lateral pinch), hand shapes (e.g. hook), hand potentials (e.g. medium wrap), object characteristics (e.g. small diameter), and opposing forces (e.g. thumb-index finger). These grasp types are systematically arranged in a tree-like fashion giving a coarse overview of how object geometry and task requirements (forces and motions) influence the grasp choice in single-handed operations by machinists working with metal parts and hand tools. When choosing a grasp type, either geometry or task requirements are predominant. For example, the grasp type used to pick up a hammer and place it in a box is different to when the task is to pound in a nail with it. On the other hand, the object geometry may be of more relevance. If the task is to file a piece of metal, the grasp type used with a small triangle file is different from that used with a large flat file, although the motions are nearly identical.

Cutkosky [1989] revised and extended the original taxonomy by adding numbers of virtual fingers to the grasp types. As substitute for this number, most grasp types of the power grasp category are termed: ”fingers surround part”. Because in a power grasp, the fingers completely envelope the object and have no independent contact areas, it is difficult to declare a number of virtual fingers. But even this extended, highly detailed taxonomy is far from covering all possible grasps humans can perform. For example, it is missing a grasp people use in writing with a pencil or in marking items with a scribe. Cutkosky [1989] had also to admit that the machinists in their study adopted numerous variations on the grasps from the taxonomy. The reasons why were ascertained to be geometric constraints, particularity of the task, personal preferences, and differences in size and strength of the hands.

2.3.2 Approach to a Taxonomy for Robot Grasping

Most of the prominent grasp taxonomies discussed in the previous section do not consider the realisation of their grasp types with a robot hand. Although authors like Cutkosky and Wright [1986] or Iberall [1997] analysed human grasping for facilitating progress in the domain of robot grasping, all of the 16, respectively 59, different grasp types they proposed cannot be realised by the majority of robot hands currently in existence. There is a need for a taxonomy for robot grasping comprising a basic set of grasp types that can be realised with most robot hands. A commonly used taxonomy provides a basis to port developmental progress in robot grasping between robot hand systems and to evaluate grasp potentials of different robot hands.

Therefore, we propose a *Taxonomy for Robot Grasping* distinguishing grasp types that are realisable by most anthropomorphic robot hands. As minimum requirements, the robot hand in use must have at least three fingers and 4 DOF like, for example, the Barrett Hand [Townsend, 2000]. The taxonomy is inspired by different studies of human grasping as those presented in the previous section. It consists of four grasp types t: *all finger precision* t_1, *two finger precision* t_2, *power* t_3, and *two finger pinch* t_4. These grasp types realise the major hand potentials and oppositions of human grasping as known from [Napier, 1956] (precision and power) and from [Iberall, 1987] (pad opposition, palm opposition, side opposition). The relations between the grasp types we propose and those of most prominent taxonomies are listed in Table 2.1.

The proposed grasp taxonomy exhibits the following features:

- small number of general grasp types;
- major potentials of human grasping can be realised (precision, power, pinch);

Research	Name of the grasp type			
Robot Grasping	all finger precision t_1	two finger precision t_2	power t_3	two finger pinch t_4
Schlesinger (1919)	palmar prehension	tip precision	spherical grasp	lateral prehension
Napier (1956)	precision grasp	precision grasp	power grasp	–
Iberall/Lyons (1984)	basic precision-b	basic precision-a	modified power	basic prec./power-d
Arbib et al. (1985)	thumb, finger 2-3	thumb, finger 2	palm, Finger 2-3	finger 2, finger 3
Lyons (1985)	precision grasp	precision grasp	encompass grasp	lateral grasp
Cutkosky (1986)	thumb-two finger	thumb-index finger	prehensile power	lateral pinch
Iberall (1987)	pad opposition	pad opposition	palm opposition	side opposition

Table 2.1: Comparison of the proposed grasp types t_1 to t_4 with established grasp taxonomies. A three-fingered hand is assumed for taxonomies that distinguish between grasps with three, four, and five fingers (Cutkosky, and Arbib et al.). The row of Arbib et al. indicates the mappings to their ”virtual finger 1” and “virtual finger 2”.

- applicable to most robot hands possessing at least three fingers;
- descriptive naming of grasp types.

While still incorporating most of the dexterity of human grasping, the small number of varying grasp types keeps the effort for development and implementation into a robot hand system low. The names of the grasp types reflect the hand potentials (precision, power, pinch) and implies the number of fingers used (two, all) except for the power grasp. Because the taxonomy is appropriate for robot hands with different numbers of fingers, the term ”all” is used in the all finger precision grasp. In the case of the power grasp, also all fingers of the robot hand are used to enclose the target object.

The four grasp types provide different precision and power capabilities. The applicability of each grasp type depends on intrinsic properties of the target object such as size, shape, and weight as well as on task and environment requirements.

- **All Finger Precision** t_1

  For grasping a broad range of middle-sized objects, this precision grasp is used in which all fingers of the robot hand are involved. The main characteristic of a precision grasp is that the target object is touched with the pad of a fingertip or near this area. With their fingertips, humans are able to place contacts and exert forces on an object with high precision because the fingertips are endowed with the highest concentration of tactile sensors. Robot hands try to copy this characteristic. Most robot hands that do not possess fingertip sensors are at least equipped with fingertips more suitable to grasp objects than other hand parts.

- **Two Finger Precision** t_2

  In contrast to the all finger precision grasp, this grasp type is only executed by the thumb and the index finger. With a thumb opposing only one finger, the gap between the grasping fingertips can be very small so that a high precision grasp type can be realised. This is needed when small or thin objects have to be grasped.

- **Power** t_3

  For grasping large, roundish, or heavy objects the grasp type to be applied has to meet high demands of stability. In a power grasp, the fingers envelop the target object. Besides the fingertips, other phalanges touch the object, and ideally, also, the palm provides opposing forces. Thus, many more contact points than in precision grasps are established. This leads to a powerful grasp.

- **Two Finger Pinch** t_4

  This grasp type combines precision and power capabilities. For the contact areas, we use the same definition as Iberall et al. [1986] use for their “side opposition” grasp: ”either between the thumb pad and the side of the index finger, or else between the sides of the fingers”. The first option is used, for instance, when holding a key to open a door. The second alternative is not considered in most taxonomies, but it is applied by humans when small objects need to be removed from tight places such as grasping coins from pockets. In robot grasping, it can be useful when obstacles constrain the environment.

The proposed grasp taxonomy provides the basic grasp types t_1 to t_4 that a robot hand should be able to carry out. This basic set is extendable by additional grasp types. To realise a grasp type t in a robot hand setup, a grasp g that fulfils the requirements described above has to be determined. The development of standard grasps for every grasp type is discussed in Section 3.2.3.

3 Development of Grasps for Robot Hands

The biologically motivated grasp strategy developed in Chapter 2 can be implemented in different real robot hand setups. During the development of this strategy and the grasp types proposed in Section 2.3.2, we had a robot setup at our disposal that included a three-fingered artificial hand. With the availability of a second robot hand setup including a five-fingered hand, we got the opportunity to prove the portability of our methods developed previously.

Both robot hand setups are described in Section 3.1, before we discuss the differences in the implementations of the portable grasp strategy. The realisation of the grasp types in both setups is examined in Section 3.2. Besides the four grasp types of our taxonomy for robot grasping, we realised a fifth type for grasping a particular object. The developmental process of grasps leads to general development rules we propose in Section 3.2.4.

3.1 Robot Hand Setups

Existing robot hand setups provide different functionalities to different extents. There are setups consisting of the robot hand only, up to systems providing multi-modal interaction capabilities [Steil et al., 2004]. To grasp objects arranged at different locations within a specified workspace, a robot hand setup at a minimum must be endowed with a device for hand positioning. For this purpose, usually a robot arm with 6 DOF is used.

The core piece of a robot hand setup is the robot hand. 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. Examples are: the Shadow Hand [Shadow Robot Company, 2006], the DLR Hand II [Butterfass et al., 2004], the Robonaut Hand [Lovchik and Diftler, 1999], the Ultralight Hand [Schulz et al., 2001], the Fluidic Hand [Schulz et al., 2004], and the GIFU Hand III [Mouri et al., 2002]. Other well known robot hands are: the Barrett Hand [Townsend, 2000], the DIST Hand [Caffaz and Cannata, 1998], the Belgrad/USC Hand [Bekey et al., 1990], the Utah/MIT Hand [Jacobsen et al., 1986], and the Stanford/JPL Hand [Salisbury, 1985].

These hands differ in size, materials, mechanical structure, number of fingers, number of joints, number of controlled degrees of freedom (DOF), actuation type and location, and sensory equipment. The number of fingers range from three to five and the number of DOF from four, in the case of the Barrett Hand, to 20, in the case of the Shadow Hand.

We use two robot hands to develop, optimise, and evaluate human-based grasps for robot grasping. When we were developing the first grasps, only the three-fingered TUM Hand (Technical University of Munich) [Menzel et al., 1993] was available at our lab. This hand is mounted on a 6-DOF PUMA robot arm and integrated into the GRAVIS system (Gestural Recognition Active VIsion System) [McGuire et al., 2002], developed within the framework of the Collaborative Research Centre (CRC) 360 “Situated Artificial Communicators” [Fink et al., 2006]. This robot system combines visual attention and gestural instruction with an intelligent interface for speech recognition and linguistic interpretation.

When we re-created the GRAVIS system, we wished to improve the hardware accuracy and to achieve the ability to investigate more humanlike grasping and manipulation of objects. More

up-to-date hardware was available and employed. The main components of our second setup are the 20-DOF Shadow Dextrous Hand [Shadow Robot Company, 2006] and the 7-DOF Mitsubishi PA-10 robot arm. The Shadow Hand is used as the end-effector of the PA-10.

The availability of these two robot hand setups that can be applied to the problems of mimicking humanoid arms provides great opportunities for the development and the investigation of robot grasping. But also great challenges are presented over the issues of platform independence and portability because the hands we use differ substantially in all characteristics mentioned previously.

3.1.1 The TUM Hand

The robot hand developed at the Technical University of Munich (TUM) consists of three identical, approximately human-sized fingers [Pfeiffer, 1996]. It has 9 DOF and is driven by an oil hydraulics system. To control the 3 DOF of each finger, motor pistons located at an external hydraulic base station drive three finger pistons by pressing oil through a hydraulic tube with a length of 2.5 m.

The fingertips are equipped with custom built force sensors to provide force feedback for control and the evaluation of grasps [Jockusch, 2000]. The hand is further equipped with a wrist camera to obtain local visual feedback during the grasping phase. The fingers do not possess sensors for directly measuring the joint angles. The only source of information about the state of the finger joints are oil pressure sensors and potentiometers for each motor piston. All sensors are located at the hydraulic base station.

In the original form of the TUM Hand shown in Figure 3.1a, the three fingers were equidistant around a circular mounting. In this configuration, the hand was only able to carry out a single type of grasp which Cutkosky and Wright [1986] termed “the precision tripod grasp”. When grasping an object, each finger provided opposing forces to both the other fingers. Thus, this grasp was a “3-virtual-finger grasp” according to the taxonomy of Arbib et al. [1985].

In the case of a human hand, there is a coupling between the index and the middle finger and they close in similar directions. The thumb is not just another finger, but can exhibit different opposing postures and can exert larger forces. With additional support from the palm, the human hand achieves the dexterity to realise much more grasp types.

Reconfiguration of the TUM Hand

In its original configuration, the TUM Hand was only a three-fingered gripper, but not an anthropomorphic robot hand, because it did not possess a finger that performed as an opposable thumb. Based on the mechanical characteristics of the fingers, we developed a palm and rearranged the three fingers so that one finger now acts as a *thumb* (see Figures 3.1b to 3.1d). We define the two other fingers to be the *index finger* and the *middle finger*. With this reconfigured, anthropomorphic TUM Hand, it is possible to realise the grasp types defined in Section 2.3.2.

Posture Control of the TUM Hand

Because the hand does not provide joint angle sensors, posture control has to be achieved indirectly, relying on the piston potentiometers and the pressure sensors, located at the base station

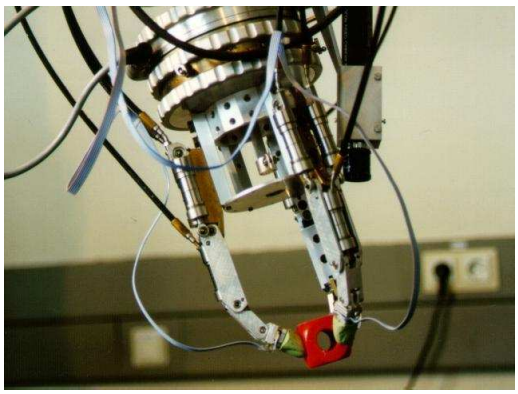

(a) Original configuration.

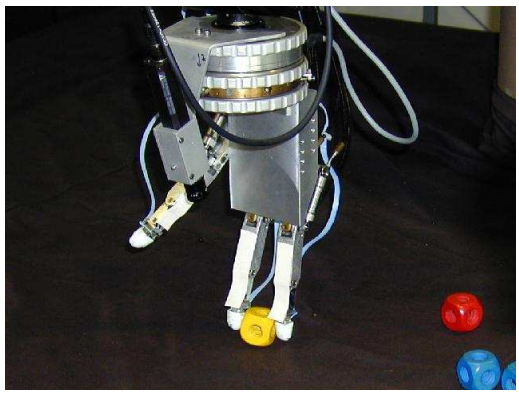

(b) Pinch grasp (anthropomorphic TUM Hand).

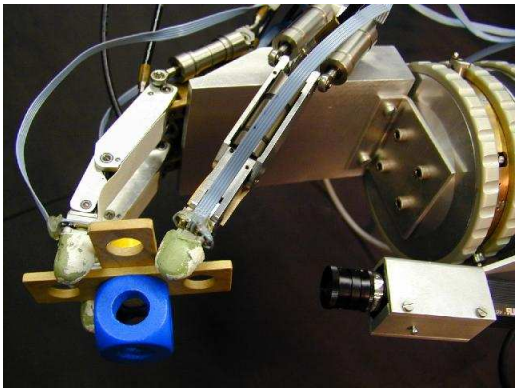

(c) Precision grasp (anthropomorphic TUM Hand).

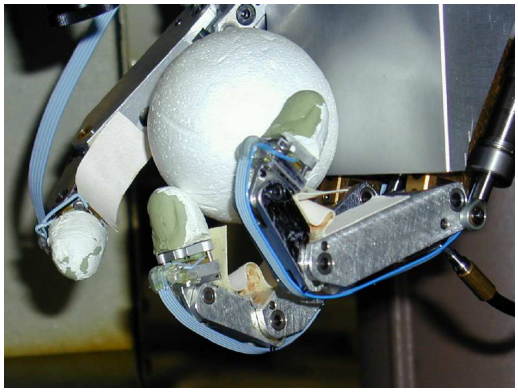

(d) Power grasp (anthropomorphic TUM Hand).

Figure 3.1: After developing a palm and rearranging the three fingers, the TUM Hand achieves anthropomorphic capabilities.

of the hydraulics system. To achieve this, we convert joint angles to piston potentiometer values by applying a fixed transform which was determined experimentally and independently for each joint.

These computed potentiometer values serve as targets for PID controllers which actuated the finger joints to move them to the desired posture. Additionally, we have to cope with hysteresis and non-linearities due to the long distance of 2.5 m 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 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.

3.1.2 The Shadow Hand

The Shadow Hand is a product of the Shadow Robot Company and is available as a prototype since end of 2004 [Shadow Robot Company, 2006]. It includes five fingers arranged at a palm and 20 actively controllable joints (20 DOF). Figure 3.2 summaries the finger kinematics. It shows a photograph of the human-like sized real hand and its kinematical model. Joint axes are visualised 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 Mallon et al. [1991]. To endow the thumb with a similar dexterity

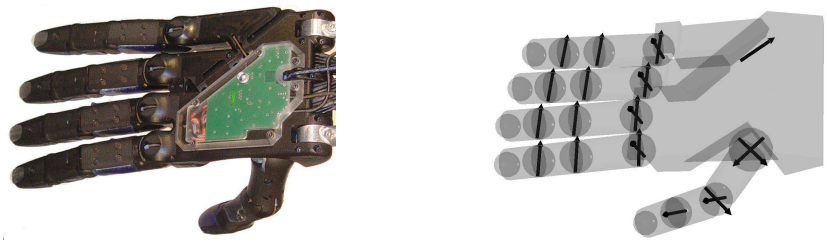

Figure 3.2: Real Shadow Hand (left) compared to its kinematic model (right). Joint axes are visualised as black arrows.

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 Figure 3.2), which allow a flexion / extension as well as abduction / adduction movement of the whole hand. Altogether, the hand includes 24 joints, 20 of them actively controllable.

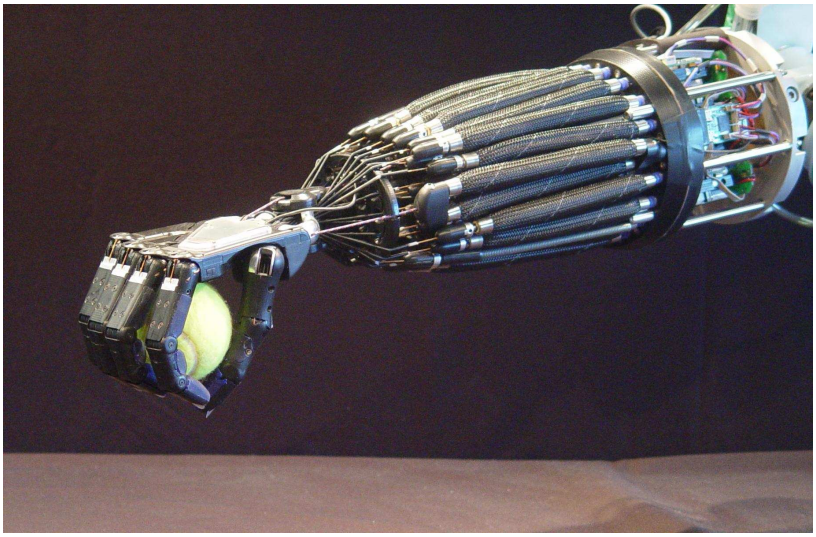

Figure 3.3: Shadow Hand grasping a tennis ball by utilising a power grasp implying palm contact.

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 Figure 3.3) and the joints are actuated by means of tendons routed through the wrist and hand. The air flow in and out of the muscles is controlled by 80 miniature solenoid on-off valves (one inlet and outlet valve for each muscle). The most important advantage of artifical muscles is their inherent and variable compliance allowing safe operation, especially in direct contact and in interaction with humans. But there are some drawbacks as well: due to friction between the tendons and their tubing as well as between the braid and the inner rubber tube of the muscles themselves, the relation between pressure, contraction, and force exhibits hysteresis and other nonlinear behaviour which complicates the control [Medrano-Cerda et al., 1995].

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 variable compliance introduced by the antagonistic pneumatic actuation, the flexibility of the

fingertips and contact areas, and the high dexterity achieved at a human size allows the Shadow Dextrous Hand to grasp everyday objects of all kind of sizes, forms, and materials having a mass of up to 5 kg.

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 approximately 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 position using Hall-effect sensors, as well as muscle air pressure. The hand system incorporates a total of 11 PIC microcontrollers, 5 of them located in the distal phalanges, one in the palm, and the remaining 5 at the base of the hand system. All controllers are connected to a single 1 Mbit CAN bus, connecting also to the controlling PC. The average total power consumption is about 25 W under normal operation (not including the external compressed air supply).

Control of the Shadow Hand

While pneumatic actuators are well known and their control has been studied mainly for single McKibben muscles [Hildebrandt et al., 2002] or actuated arms [Tonietti and Bicchi, 2002], 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 modelling scheme like that proposed in Granosik and Borenstein [2004] becomes difficult to apply and would have to be carefully adapted to each muscle. The tight packaging of the actuator muscles in the forearm forms a second source of friction and additionally causes a coupling between DOF which in rare cases even can cause single muscles to get temporarily stuck together. On the other hand, for applying the grasp strategy proposed in Section 2.3.2 a sophisticated joint control is not required as long as the pre-grasp and the target grasp postures can be approximately actuated. Therefore, we decided to base our joint controller on a relatively simple mixing scheme designed to control position together with stiffness and having only moderate parameter tuning complexity. The working principle of the mixing controller is summarised in Figure 3.4.

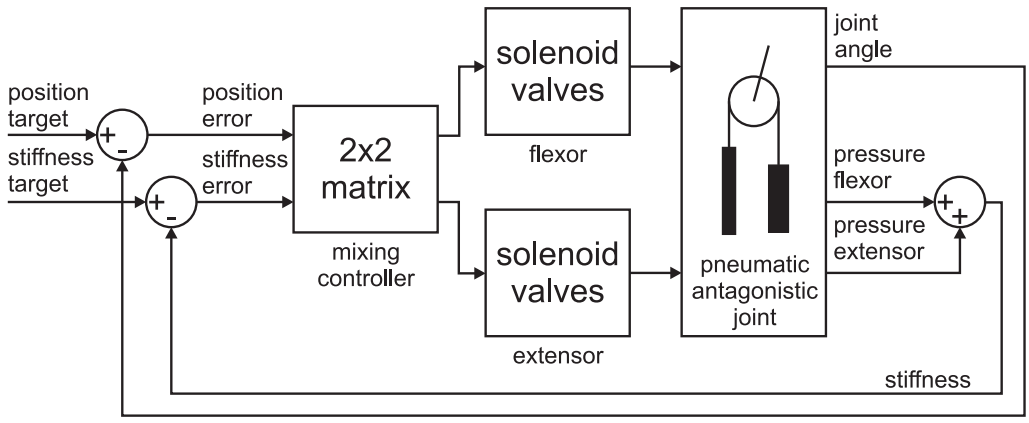

Figure 3.4: Schematic view of the mixing controller allowing simultaneous stiffness and position control.

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. One solution presented in [Vanderborght et al., 2004; Schröder et al., 2003] employs a pressure change Δp computed from the joint error and feeds this as $p_{\text{work}} \pm \Delta p$ into the pressure controllers of the agonist and antagonist respectively, where p_{work} is the average working pressure of the muscles.

The key idea of the controller is the combination of two control variables, joint position and joint stiffness, which is motivated by the observation that the pressure difference correlates with the joint position while the pressure sum correlates with the stiffness of the joint [Granosik and Jezierski, 1999]. 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:

$$\begin{pmatrix} t_{flex} \\ t_{ext} \end{pmatrix} = \begin{pmatrix} K_\theta & K_S \\ -K_\theta & K_S \end{pmatrix} \begin{pmatrix} \Delta\theta \\ \Delta S \end{pmatrix} \tag{3.1}$$

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 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.

3.1.3 Implementation of the Grasp Strategy

Before the grasp process of the strategy proposed in Section 2.3.2 can be carried out, the grasp to be applied has to be chosen. The development and the implementation of the standard grasps for each grasp type t is described in Section 3.2. Additionally, optimised versions of object-specific grasps are available in both robot hand setups after applying the optimisation strategy described in Section 6.3. In both setups, each of these hand-dependent grasps can be chosen irrespective of the target object.

TUM Hand Setup

The implementations of our grasp strategy differ between the two robot hand setups. The multimodal interaction setup integrating the TUM Hand provides the possibility of interactively choosing the target from objects arbitrarily arranged in the workspace. The grasp type can be selected by observation and 3D-identification of human hand postures. The realisation of these capabilities is out of the scope of this thesis but is described in detail in Ritter et al. [2003]. After choosing the grasp to be applied, the grasp strategy is realised by the following steps:

1. *Approach/Pre-grasp phase*: The target object is approached after the object's 3D-coordinates are resolved by a stereo vision system to an accuracy of about 3 cm. A visually guided fine-positioning step based on the wrist camera improves the error to about 1 mm and centres the manipulator above the object. Then the wrist camera is used to analyse the horizontal orientation of the object which is the final parameter needed to determine the transformation between the hand frame and the object-zero frame (defined in Section 3.2.2). This allows the system to compute the grasp type specific pre-grasp position which is subsequently approached. Concurrently the pre-grasp is actuated.

2. *Placing phase*: After the pre-grasp position is reached and the hand is in pre-grasp posture, the grasp position is approached.

3. *Grasp closure phase*: When the grasp position is reached, the grasp is applied.

4. *Stabilisation phase*: The fingertip force sensors, allow the evaluation of grasp stabilisation, i.e. determining if grasp type specific force thresholds are reached.

5. *Lift-off phase*: When the grasping fingers exert sufficient forces on the target object, the hand is moved back towards the pre-grasp position.

Shadow Hand Setup

At the time when we implemented the grasp strategy on the **Shadow Hand setup**, no vision system was available that could guide the positioning of the robot hand in relation to the target object. Thus, we define position and orientation of the object zero-frame and supply them to the robot system. Before the grasp strategy is applied, the target object has to be positioned according to this fixed frame. Another limitation that we had to cope with was that the first force sensor prototypes of the Shadow Hand showed strong sensitivity fluctuations. That complicated the determination of force thresholds for stable grasping to an extent that we chose to use no tactile feedback in the implementation of our grasp strategy. These differences compared to the TUM Hand setup changed the implementation of the grasp strategy as follows:

1. *Approach/Pre-grasp phase*: The predefined, grasp type specific pre-grasp position is approached. Concurrently the pre-grasp is actuated.

2. *Placing phase*: After the pre-grasp position is reached and the hand is in pre-grasp posture, the grasp position is approached.

3. *Grasp closure phase*: When the grasp position is reached, the grasp is applied.

4. *Stabilisation phase*: For the grasp closure phase and the stabilisation phase, a time of five seconds is assumed to be sufficient in the case of all grasp types except for the two finger precision t_3. When grasping small objects, the two finger precision grasp nearly reaches its target grasp posture. The parameterisation of the Shadow Hand controller (see Section 3.1.2) produces slower finger movements when reaching the targets. Thus, for grasp type t_3 a time of ten seconds is provided.

5. *Lift-off phase*: When the time for grasp closure and stabilisation has elapsed, the hand is moved back towards the pre-grasp position.

3.2 Development of Grasps

To realise the grasp types t_1 to t_4 from the taxonomy for robot grasping in a robot hand setup, *standard grasps* g_1 to g_4 have to be defined that meet the requirements proposed in Section 2.3.2. For this purpose, appropriate parameterisations of g_1 to g_4 are determined in a developmental process. As introduced in Section 2.2.4, a grasp includes two joint angle configurations: pre-grasp and target grasp. Additionally, the position $\boldsymbol{p}$, the orientation $\boldsymbol{o}$, and the approach distance d have to be identified.

The determination of these parameters is the subject of this section. After identifying two parameters that are hand-independent, the developmental process of realising grasp types for the two robot hands we use is discussed, and the parameterisations of the standard grasps are presented. Additionally, a fifth grasp type, the *three finger special* t_5, is defined and the development of its standard grasp g_5 is described. The grasp type t_5 extends our taxonomy for robot grasping and is used for grasping a particular benchmark object, the toy propeller. Sections 4.1.1 and 4.2.2 provide more discussion on this object.

3.2.1 Hand-Independent Parameters

Two parameters of the standard grasps can be determined independently of the robot hand used. One of these parameters is the *approach distance* d. This distance distinguishes the pre-grasp position from the grasp position. In the task of grasping objects lying on a desktop from above, the hand movement in the placing phase is directed perpendicular towards the desktop surface. The only constraint on the approach distance d is that it has to be large enough to avoid a contact of the hand with the target object during the approach/pre-grasp phase. We found that for both robot hands and for every grasp type an approach distance of $d = 8\,cm$ is sufficient. Thus, besides being hand-independent, the parameter d is equal for every standard grasp.

The second parameter that does not depend on the robot hand under investigation is the azimuth angle.

The Azimuth Angle α

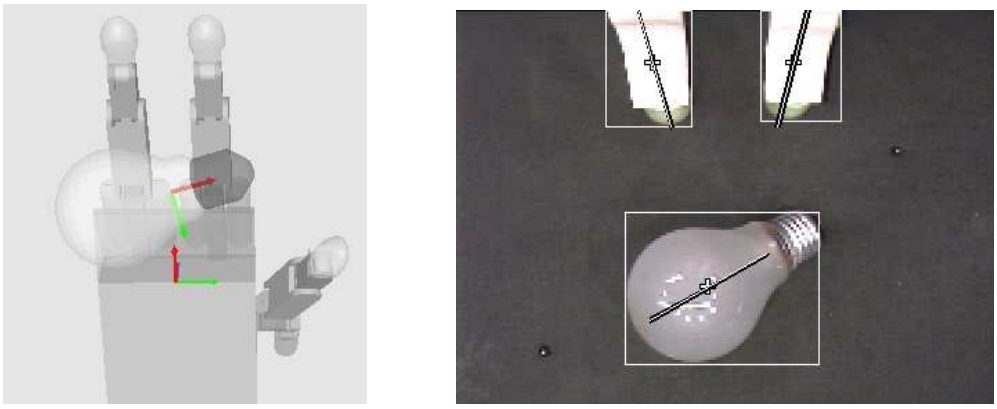

(a) Red arrows of respective frames denote hand and object orientation.

(b) The orientation of the object is analysed in the wrist camera picture during the approach phase.

Figure 3.5: Orientation of the hand and the target object.

The *hand azimuth* is defined as the projection of the orientation vector of the hand into the horizontal plane, i.e. the plane of the desktop surface. The hand's orientation vector is parallel to the palm and is directed towards the outstretched fingers (see Figure 3.5a). The *azimuth angle* α is the angle between the hand azimuth and the horizontal orientation of the target object, which in the case of the TUM Hand is known from visual object recognition (see Figure 3.5b). The horizontal orientation vector of an object is defined to be parallel to the longest side of the horizontal projection of the object. If this projection is a circle (for instance, in the case of balls or upright standing cylinders) no distinctive horizontal orientation can be determined.

In many studies on conditions that influence the orientation of the grasping hand, objects with no distinctive horizontal orientation have been used, and studies that relate hand and object orientations are rare. Paulignan et al. [1997] showed that the orientation of the hand in relation to the object mainly depends on object position. To be more precise, the authors examined the orientation of the opposition axis (defined as the line connecting the tips of the thumb and the index finger). They used cylindrical objects placed upright on a horizontal plane and thus no horizontal object orientation could be determined. Other studies showed that not only the position of the object but also the grasping movement direction determines the hand orientation in the grasp closure phase (for example, Bennis and Roby-Brami [2002] and Roby-Brami et al. [2000]). In these experiments, cylindrical objects standing upright were used too.

Objects with decisive horizontal orientation (bars) had to be grasped in an experiment conducted by Fan et al. [2005]. The authors found that the target object orientation is a constraint on the hand/arm orientation at the time of grasping. Their results show that there is a difference of approximately 75 degrees between the horizontal orientations of the hand and the object. This difference is independent of target orientation and is unaffected even when this orientation is changed during the grasp trial.

In preliminary experiments of grasping the 21 benchmark objects with different orientations, we found that this difference, azimuth angle $\alpha = 75°$, leads to the most successful grasp trials for both of the robot hands we use. If the target object is spherical or an upright cylinder, of course α can take any value. But if an object with determinable horizontal orientation has to be grasped with a precision t_1 or t_2 grasp type or a power t_3 grasp type, an azimuth angle $\alpha = 75°$ is appropriate (compare orientation of red arrows in Figure 3.5a).

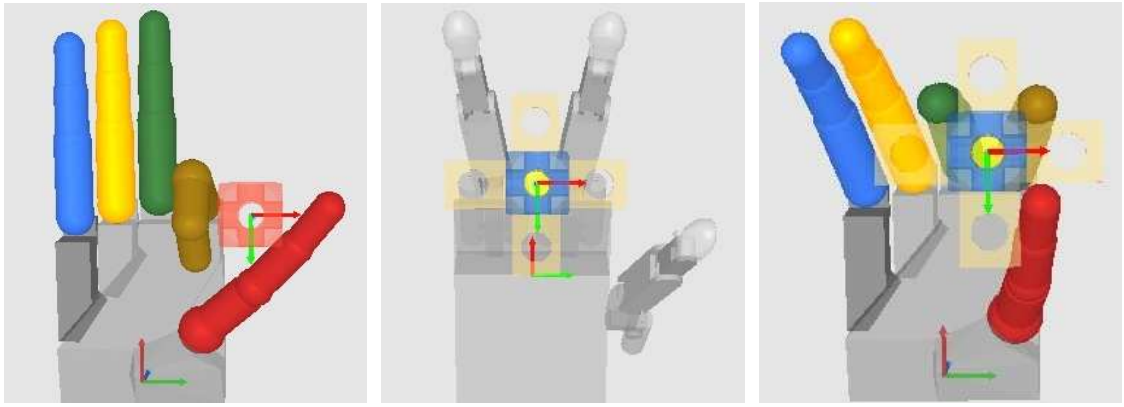

(a) Shadow Hand grasping the toy cube (two finger pinch t_4 pre-grasp posture).

(b) TUM Hand grasping the toy propeller (three finger special t_5 pre-grasp posture).

(c) Shadow Hand grasping the toy propeller (three finger special t_5 grip posture).

Figure 3.6: Simulated grasps viewed from below the desktop. Hand azimuth equals the object orientation minus 90 degrees (red arrows are perpendicular) in the cases of the two finger pinch t_4 grasp type and the three finger special t_5 grasp type.

An azimuth angle $\alpha = 90°$ is suitable for the two remaining grasp types. In the case of the two finger pinch t_4, the reason is that the side of one or both fingers is used to grasp the toy cube (which is the only object being grasped with the two finger pinch t_4), and thus the orientation of the finger side and that of the object's contact side should be equal (see Figure 3.6a). The three finger special t_5 considers the complex shape of the object (toy propeller), and the grasping fingers which are widely spread in pre-grasp posture (see Figure 3.6b) have the best chance to enclose the object (see Figure 3.6c) when the hand and the object orientations are perpendicular.

3.2.2 Coordinate Frames

The grasp position of the hand is defined by the position vector $\boldsymbol{p}$. The orientation vector $\boldsymbol{o}$ determines the orientation of the hand in both the pre-grasp position and the grasp position. For every standard grasp, the two vectors, $\boldsymbol{p}$ and $\boldsymbol{o}$, are related to the target object. Before $\boldsymbol{p}$ and $\boldsymbol{o}$ of each standard grasp can be determined, the coordinate frames related by $\boldsymbol{p}$ and $\boldsymbol{o}$ have to be identified.

Each of the standard grasps has to be suitable for grasping objects of different sizes and shapes. For grasping even small and flat objects, it is essential that the fingers close directly above the desktop surface when applying a standard grasp. Therefore, $\boldsymbol{p}$ relates to the projected point of the object's centre of mass into the plane of the desktop surface, while $\boldsymbol{o}$ relates to the horizontal projection of the object's orientation.

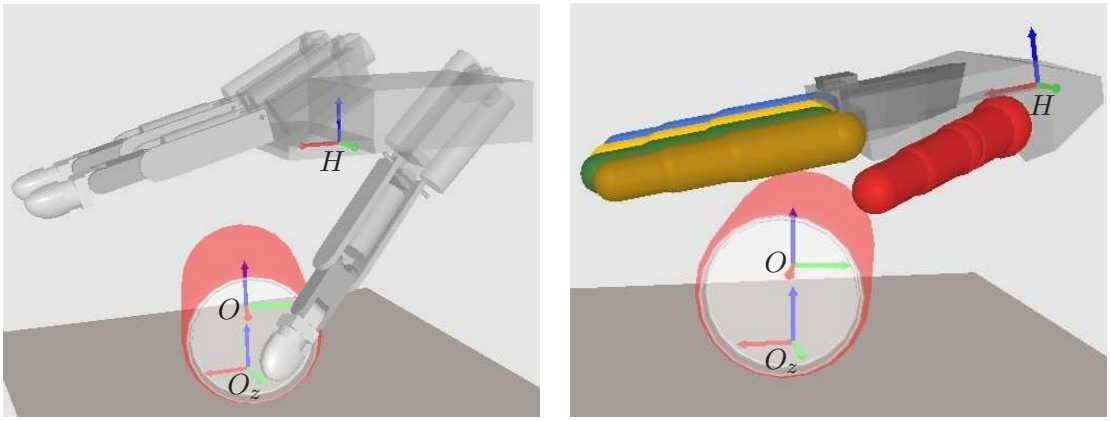

(a) Hand frame H and object zero-frame O_z have equal orientations when the palm is in parallel to the desktop.

(b) Object frame O and object zero-frame O_z are distinguished by a rotation α around z_{O_z} and a translation along it.

Figure 3.7: Illustration of the hand frame H, the object frame O, and the object-zero frame O_z. Red arrows correspond to x-axes, green arrows to y-axes, and blue arrows to z-axes.

We define two **coordinate frames** depicted in Figure 3.7:

- The *hand frame* H
  The x-axis x_H and the y-axis y_H lie in the plane of the palm, where x_H is directed towards the outstretched fingers. The z-axis z_H is directed towards the dorsal side of the palm and y_H completes a right-handed coordinate system. The origin o_H of the hand frame H can be defined for each hand arbitrarily.

- The *object zero-frame* O_z
  The origin of the target object zero-frame O_z is the perpendicular projection of the object's

centre of mass on the desktop surface. The z-axis z_{O_z} is the perpendicular of the desktop and is directed upwards. The x-axis x_{O_z} lies in the direction of the horizontal projection of the object's orientation rotated by the negative azimuth angle $-\alpha$ around the z-axis z_{O_z}. The y-axis y_{O_z} completes a right-handed coordinate system.

With this definition of coordinate frames, the position vector $\boldsymbol{p} = (x, y, z)^{-1}$ determines the position of the hand by the translations: x along x_{O_z}, y along y_{O_z}, and z along z_{O_z}. The hand is oriented according to the orientation vector $\boldsymbol{o} = (\gamma, \beta, \alpha)^{-1}$. With the palm in parallel to the desktop surface, the hand is rotated by γ (roll) around x_H and β (pitch) around the new y_H. The yaw angle equals α and is considered in the definition of the object zero-frame O_z.

3.2.3 Hand-Dependent Parameters

Without the two hand-independent parameters, d and α, the parameters to be determined for each of the five standard grasps, *all finger precision* g_1, *two finger precision* g_2, *power* g_3, *two finger pinch* g_4, and *three finger special grasp* g_5, are:

- x, y, and z of position vector $\boldsymbol{p}$,
- γ (roll) and β (pitch) of orientation vector $\boldsymbol{o}$,
- a set of joint angles identifying the pre-grasp,
- a set of joint angles identifying the grasp.

The number of joint angles depends on the hand in use. Overall, 23 parameters have to be determined in the case of the TUM Hand and 41 parameters in the case of the Shadow Hand. These values are interdependent and are strongly related to the geometry of the palm and the fingers of the robot hand used. But they can be experimentally determined for every standard grasp in a developmental process consisting of the following iterative steps:

- Rearrange position and orientation of the hand.
- Modify joint angle values of pre-grasp and target grasp.
- Execute trials of grasping different objects.

The number of iteration steps that is necessary to realise a grasp type, fulfilling the requirements described in Section 2.3.2 and leading to a satisfactory grasp success, varies. It depends on the grasp type to be realised, the hardware in use, whether or not an appropriate grasp simulator is available, and on the experience of the developer. The general development rules presented in Section 3.2.4, following the discussion of the realisation of the standard grasps, facilitate this process.

The parameters of position vector $\boldsymbol{p}$ and orientation vector $\boldsymbol{o}$ determined for the standard grasps g_1 to g_5 are listed in Table 3.1. As already described, the yaw angle equals the azimuth angle, and $\alpha = 90°$ for the grasp types t_4 and t_5, while $\alpha = 75°$ otherwise.

The values of the roll angle γ and those of the pitch angle β do not differ between the five standard grasps in the case of the Shadow Hand. In contrast, these values show more variance in the case of the TUM Hand. The reason can be found in the dexterity of the thumb. The Shadow Hand is equipped with a thumb possessing five joints thereby having two more degrees of freedom.

standard grasps	position $\boldsymbol{p}$ (in cm)			orientation $\boldsymbol{o}$ (in degrees)		
	x	y	z	γ (roll)	β (pitch)	α (yaw)
TUM-Hand						
all finger precision g_1	-6.4	-0.2	11.4	-1.8°	19°	75°
two finger precision g_2	-3.9	-1.7	11.8	-1.5°	25°	75°
power g_3	-5.4	-0.2	11.3	-1.5°	26°	75°
two finger pinch g_4	1.6	0	17.0	0°	90°	90°
three finger special g_5	-4.8	0.1	13.7	-1.5°	26°	90°
Shadow-Hand						
all finger precision g_1	-9.5	-1.2	11.1	-1.5°	10°	75°
two finger precision g_2	-9.2	-3.2	11.0	-1.5°	10°	75°
power g_3	-10.5	-0.2	10.6	-1.5°	10°	75°
two finger pinch g_4	-8.5	-6.0	8.2	-1.5°	10°	90°
three finger special g_5	-9.8	-2.4	13.0	-1.5°	10°	90°

Table 3.1: Position $\boldsymbol{p}$ and orientation $\boldsymbol{o}$ identified for all standard grasps of both hands.

The lack of dexterity in the thumb of the TUM Hand has to be compensated for by adjusting the orientation angles γ and β to achieve a thumb trajectory close to the desktop surface.

The following detailed discussion about realising the standard grasps provides more explanation of the parameters listed in Table 3.1 and a qualitative description of the pre-grasps and the target grasps. Detailed listings of the values of the joint angle configurations are presented in Section A.1.

Realising the Standard Grasps of Grasp Types t_1 to t_5

- **All Finger Precision Grasp** g_1

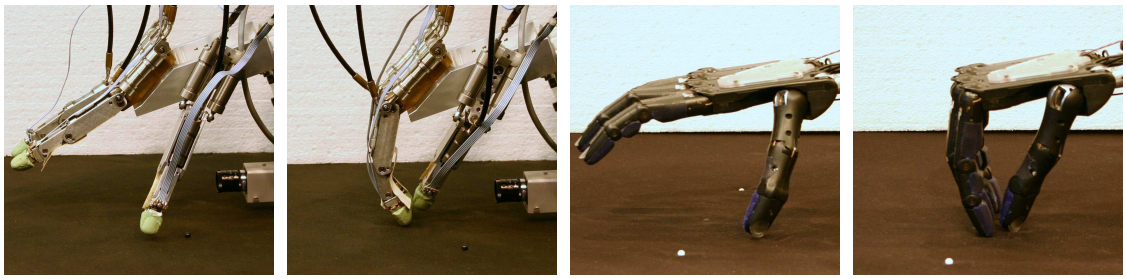

(a) Pre-grasp (TUM). (b) Target grasp (TUM). (c) Pre-grasp (Shadow). (d) Target grasp (Shadow).

Figure 3.8: Pre-grasp and target grasp postures of the *all finger precision* grasp g_1.

As the major characteristic of precision grasps is that the target object is touched with the fingertips, in the target grasp posture of the *all finger precision* grasp g_1 (see Figures 3.8b and 3.8d) the pads of the fingertips have to approach each other. But they do not touch, to allow both the detection of a successful grasp and the miss or loss of the object by analysing the values of the fingertip force sensors (if they can be used).

In pre-grasp posture, the fingers have to be opened as far as possible, to provide the ability of grasping even large objects. In the case of the TUM Hand, the fingers are completely extended (see Figure 3.8a), but the fingers of the Shadow Hand have to be slightly flexed (see Figure 3.8c) because otherwise they could stick into the desktop surface during grasp closure.

- **Two Finger Precision Grasp** g_2

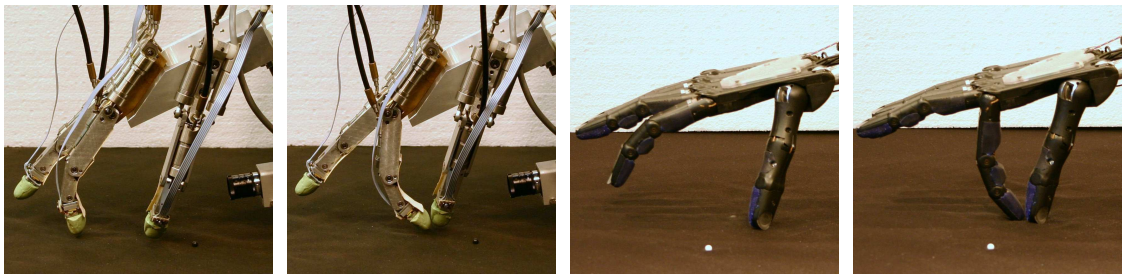

(a) Pre-grasp (TUM). (b) Target grasp (TUM). (c) Pre-grasp (Shadow). (d) Target grasp (Shadow).

Figure 3.9: Pre-grasp and target grasp postures of the *two finger precision* grasp g_2.

With the *two finger precision* grasp g_2, it has to be possible to grasp very small objects. In contrast to the all finger precision grasp, the index finger has to be flexed more, to minimise the gap between the grasping fingers and the desktop surface (see Figures 3.9b and 3.9d). This is required to realise a high precision grasp because the fingertips of the robot hands are roundish at the top and are less soft when compared with those of a human hand. There is no contact between the fingers in target grasp posture, to facilitate an analysis of the fingertip sensors according to the grasp success.

Again, the pre-grasp postures (see Figures 3.9a and 3.9c) have to prevent the fingers from sticking in the desktop surface during the subsequent grasp closure phase. To prepare the flexion of the index finger in target grasp posture, a moderate flexion in pre-grasp posture is required, even in the case of the TUM Hand.

- **Power Grasp** g_3

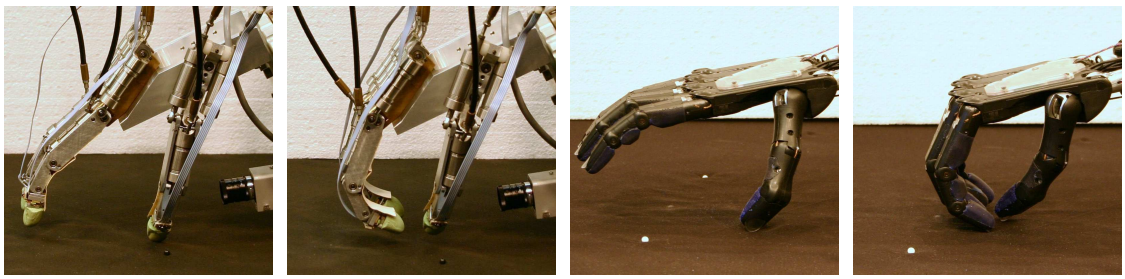

(a) Pre-grasp (TUM). (b) Target grasp (TUM). (c) Pre-grasp (Shadow). (d) Target grasp (Shadow).

Figure 3.10: Pre-grasp and target grasp postures of the *power* grasp g_3.

In the target grasp posture of a *power* grasp g_3 (see Figures 3.10b and 3.10d) the fingers have to adopt a more curved form than in the target grasp posture of a precision grasp. To enclose the target object lying on the desktop, with the fingers being more curved, the height (z) of the hand has to be a bit less than in all other grasps.

In order not to stick in the desktop surface, the fingers of the TUM Hand have to be flexed by a small amount in the pre-grasp posture (see Figure 3.10a). The finger pre-grasp postures of the Shadow Hand equal those of the all finger precision grasp, expect the thumb is more flexed (see Figure 3.10c).

- **Two Finger Pinch Grasp** g_4

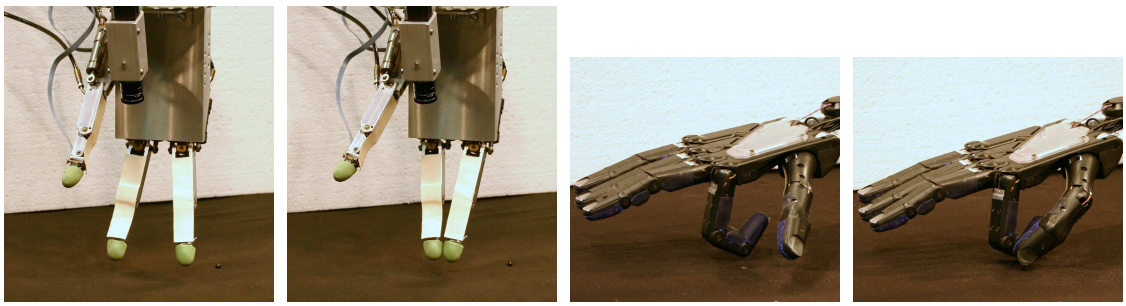

(a) Pre-grasp (TUM). (b) Target grasp (TUM). (c) Pre-grasp (Shadow). (d) Target grasp (Shadow).

Figure 3.11: Pre-grasp and target grasp postures of the *two finger pinch* grasp g_4.

There are large differences between the *two finger pinch* grasps g_4 developed for the two robot hands. The grasp types differ in the fingers used and in position $\boldsymbol{p}$ and orientation $\boldsymbol{o}$. The principle reason is that in the case of the TUM hand only a pinch grasp between the index and the middle finger is realisable. Grasping with the thumb implies the sides of all phalanges of the index finger must be used as opposing faces. But the metallic sides of the proximal and the middle phalanges of the TUM fingers are not appropriate for grasping an object, and only the sides of the fingertips can be used. But this type of pinch grasp provides an advantage when grasping in constrained environments as already discussed in Section 2.3.2.

The realisation of the pinch grasp between the index and the middle finger of the TUM Hand was uncomplicated. The orientation vector of the palm is perpendicular to the desktop surface ($\gamma = 0°$, $\beta = 90°$). In pre-grasp posture (see Figure 3.11a), the two fingers have to be abducted so that the target object can be enclosed between them. In target grasp posture (see Figure 3.11b), the fingertips have to be as close as possible to grasp even small objects. They do not touch each other for the same reason as discussed before.

In the case of the Shadow Hand, the touch areas of the phalanges extend to the sides of the fingers. Thus, the more common pinch grasp between the thumb pad and the side of the index finger can be applied. This grasp type provides larger contact areas, which facilitate grasps with higher stability when compared with the alternative pinch grasp between the sides of the fingers. As shown in Figures 3.11c and 3.11d, the thumb is mainly responsible for opening and closure of the grasp, whereas the index finger only supports it with a small adduction movement in its metacarpophalangeal joint.

- **Three Finger Special Grasp** g_5

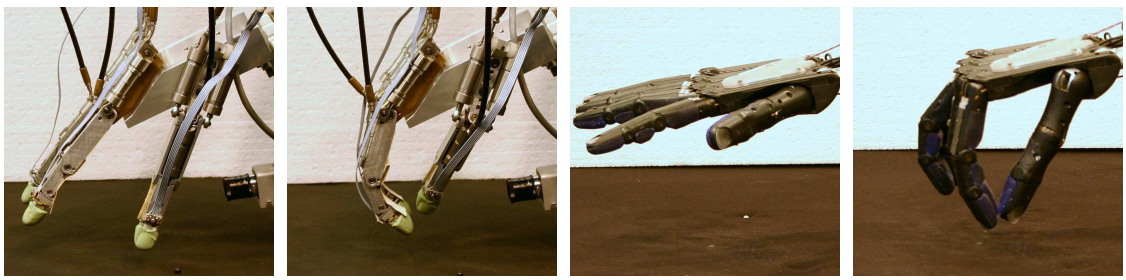

(a) Pre-grasp (TUM). (b) Target grasp (TUM). (c) Pre-grasp (Shadow). (d) Target grasp (Shadow).

Figure 3.12: Pre-grasp and target grasp postures of the *three finger special* grasp g_5.

The *three finger special* grasp g_5 is a standard grasp of a grasp type t_5 which extends our basic robot grasp taxonomy as suggested in Section 2.3.2. This grasp type has precision characteristics because the target object is touched by the pads of the fingers. It was defined for a particular benchmark object, the toy propeller (see Section 4.1), which could not be grasped with the TUM Hand by applying any of the other standard grasps. Because of the complex shape of the toy propeller, it is pushed aside by the index or the middle finger before the thumb could exert an opposing force on the object.

The three finger special grasp avoids this circumstance by an abduction of the index and middle finger in pre-grasp posture (see Figures 3.12a and 3.12c). This leads to the enclosure of a propeller blade during grasp closure and reduces the displacement of the object. The target grasp posture (see Figures 3.12b and 3.12d) leads to contact points between the blades of the propeller during grasp closure. To reach these points with the tips of the three fingers, the distance (z) of the hand to the desktop has to be larger than in the case of all other standard grasps.

3.2.4 General Development Rules

Since the TUM Hand was available for this study before obtaining the Shadow Hand, the grasp types were primarily developed for grasping with the TUM Hand and ported to the Shadow Hand afterwards. Additionally, the grasp simulation was not available at the time when the grasp types were developed, but the simulator has been useful for their adaption in the case of the Shadow Hand. By determining the 25 (two hand-independent plus 23 hand-dependent) parameters of each grasp type for the TUM Hand in the developmental process described, a lot of experience was gained that facilitated this tuning phase in the case of the Shadow Hand.

Some general rules summing up these experiences and largely facilitating this process can be formulated. The position $\boldsymbol{p}$ and the orientation $\boldsymbol{o}$ have to be adjusted such that the centre point of the grasping fingers is close to the object's centre 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 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 constraints, it is fairly easy to develop suitable realisations of the grasp types.

4 Benchmark and First Evaluation

For evaluating grasp types realised in a robot hand setup and comparing the capabilities of different robot hands and different grasping strategies, we define a benchmark system. This includes a diversified repertoire of everyday objects and defines a benchmark test. It is described in Section 4.1 after reviewing several grasp assessment tests existing in different research fields.

A first evaluation of the basic grasp types with the TUM Hand utilising the benchmark system proposed is presented in Section 4.2. The results of the experiment defined by the benchmark test are discussed and are used to choose the most suitable grasp type for each particular benchmark object.

4.1 Benchmark System for Robot Grasping

To evaluate grasping of everyday objects and to compare the performance of different robot hands, an assessment test is required. We have found hardly any evaluation tests for robot grasping and there exists no standardised benchmark. Even in the fields of hand surgery, rehabilitation, and the design of prosthetic devices standardised methods are rare, although many assessment tests for the purpose of evaluating pathologic and prosthetic hand function exist.

In a detailed review of existing hand assessment procedures, Light et al. [1999] conclude that there is no conformity in these protocols and most of them rely on subjective evaluations. These tests are commonly used on a comparative basis, to enable surgeons and therapists to evaluate the rehabilitation of an injured or dysfunctional hand. A subjective evaluation of the transition in hand impairment during and after treatment is important, but these kinds of tests cannot provide a basis for a standardised and objective procedure with quantified requirements.

A few approaches to standardised hand function tests for pathologic and prosthetic hands exist. In these tests, the term ”hand function” is not used uniformly, but generally it encompasses the ability to perform tasks encountered during everyday living [Light et al., 1999]. A hand function test, for comparing the results of reconstructive hand surgery, is proposed by Sollerman and Ejeskar [1995]. It consists of 20 activities of daily living. For each of these activities one, two, or three out of seven different grasp types are allowed. Scoring is based on time margins but also on subjective appraisal of the grasp type used. A procedure for evaluating pathologic and prosthetic hand function that results in an ”index of functionality” by only evaluating the time needed is proposed by Light et al. [2002]. 12 abstract objects and 14 everyday objects have to be grasped and manipulated. Although one or more out of six grasp types are allowed for each task, the type actually chosen by the subject is not evaluated.

In contrast to existing assessment tests of pathologic and prosthetic hand function, our benchmark has been devised for evaluating grasping with a robot hand and not for the ability of the hand to perform everyday tasks. Such tasks require a lot of different hand and arm movements that are out of the scope of this thesis. Nevertheless, our benchmark complies with the main requirements Light et al. [1999] impose on standardised assessments:

- The test must cover all possible grasp types.
- No subjective opinion should be used during the evaluation.

- The current technology should not affect the determination of the test procedures.

A first attempt to promote the definition and use of benchmarks in robotic manipulation and grasping is done by the EUropean RObotics Network [EURON, 2006]. The contributers propose three different benchmark tests [Pobil, 2006]. In the first test (i) both regularly and non-regularly shaped objects have to be grasped and re-grasped. Test (ii) consists of a pick and place operation. After lifting the object, it has to be accelerated in different directions, and the maximum acceleration is taken as the evaluation measure. The pick and place tasks in test (iii) have to be performed as quickly as possible and the measures should be time and statistics of failures.

A first experimental protocol for a benchmark based on test (ii) has been proposed by Morales [2006]. The goal of this experiment is to test different procedures for visually-guided grasping. By only using visual information of 67 planar objects, a set of three-finger grasps are determined and executed [Morales, 2003]. Two of these grasps have been evaluated for maximum acceleration to illustrate test (ii) [Morales, 2006]. This test is only appropriate if the target objects are highly likely to be grasped and lifted successfully. Then the value of maximum acceleration can be a useful quality measure.

Everyday objects with different sizes and shapes, however, are difficult to grasp with a robot hand. We propose a benchmark test that leads to a more suitable evaluation of the quality of a grasp: executing several grasp trials on each object, and taking the number of successful trials as evaluation measure. The time needed to execute the grasp is not considered in our grasp optimisation strategy proposed in Section 6.3. Thus, our benchmark test is also not related to test (iii). At best it matches the EURON test (i) for manipulation and grasping, i.e. the ”grasp and re-grasp test”.

The benchmark test we propose is presented in Section 4.1.3. It is applied to a set of 21 everyday objects. In Section 4.1.1 the various aspects considered in choosing the benchmark objects are discussed. Note that typically anybody interested in robot grasping is familiar with the chosen objects and can gather similar objects as required. If it is not possible to obtain an exact model of a benchmark object, it is necessary to have characteristics that allow the selection or construction of a similar object that poses equal challenges for the grasp process. To this end, we provide all of the characteristics that are required in Section 4.1.2. In addition, we use these characteristics to develop models of the benchmark objects in our grasp simulator introduced in Section 6.1.

4.1.1 Benchmark Objects

To create a benchmark system for the evaluation of grasping everyday objects, a diversified repertoire of objects was used. When applying the benchmark system, it has to evaluate general grasping capabilities and must not be restricted to a specific domain. Therefore, objects were chosen that humans use in a variety of everyday situations found in many different scopes of application: office equipment, toys, toiletries, sports goods, crockery, housewares, electronic equipment, food and tools. It is also valuable if the objects differ in the characteristics that affect the grasping process and lead to potentially different classes of grasp difficulty. Thus, the objects have to differ in size, shape, mass, surface condition, compliance, centre of mass in relation to the geometric centre, and plasticity. Based on these considerations 21 everyday-objects were chosen as benchmark objects which are pictured in Figure 4.1.

For categorising these objects, we distinguish between the difficulties the objects can impose on the grasp process. This process can be interpreted as a formation of the grasping fingers and is often treated as a shape matching problem (see, for example, Li and Pollard [2005]). Both the shape and the plasticity of the objects have to be considered. This leads to a categorisation divided into three main classes: 1) simple geometric shaped objects; 2) complex geometric shaped objects;

Figure 4.1: The real benchmark objects and their 3D-models in simulation. After the first evaluation of the standard grasps applied with the TUM Hand, the 21 objects were subdivided into five categories with increasing grasp difficulty (see Section 4.2).

3) form-variable objects. To provide diversity of sampled objects, the first category considers the most common simple geometric shapes, i.e.: spherical, cubic, cuboid, cylindrical and conical shapes. The selected 21 objects are categorised according to shape and formability as follows:[1]

1. **Simple geometric shaped objects:**
   - sphere: tennis ball (6), golf ball (13);
   - cube: toy cube (3);
   - cuboid: matchbox (14), eraser (19);
   - cylinder: can (4), board marker (11), tea light (12), pencil (21);
   - cone: cup (10).
2. **Complex geometric shaped objects:**
   adhesive tape (1), toy propeller (2), sharpener (8), remote control (9),
   light bulb (15), voltage tester (18).
3. **Form-variable objects:**
   tissue pack (5), paper ball (7), chocolate bar (16), folding rule (17), bunch of keys (20).

Two objects, the toy cube and toy propeller, are components of an assembly scenario investigated in the framework of the long-term Collaborative Research Centre (CRC 360) “Situated Artificial Communicators” funded by the German Research Society [Fink et al., 2006].

[1] The bracketed numbers correspond to the imaged objects in Figure 4.1.

4.1.2 Object Characteristics

The characteristics required to simulate an object are the mass, the dimensions, the offset of the centre of mass from the geometric centre in the longitudinal direction, and the two coefficients of friction (static and dynamic).

no.	object	mass	dimensions (mm)			offset	frictional coefficients	
		(g)	length	width	height	of CM	static	dynamic
1	adhesive tape	24.6	61.7	53.4	26.0	-	0.39	0.32
2	toy propeller	25.9	91.2	91.2	51.4	-	0.33	0.18
3	toy cube	12.2	31.0	31.0	31.0	-	0.31	0.14
4	can	368.0	65.8	65.8	115.6	-	0.23	0.18
5	tissue pack	25.7	107.0	53.0	25.0	-	0.35	0.32
6	tennis ball	58.4	66.0	66.0	66.0	-	0.79	0.74
7	paper ball	5.0	68.0	58.0	46.0	-	0.75	0.71
8	sharpener	24.4	36.6	36.6	57.0	-	0.44	0.42
9	remote control	145.9	165.0	50.0	40.0	11.0	0.4	0.37
10	cup	236.4	107.5	85.0	65.0	10.0	0.4	0.4
11	board marker	16.8	140.0	17.8	17.8	2.0	0.44	0.4
12	tea light	13.5	38.0	38.0	16.5	-	0.3	0.25
13	golf ball	45.7	42.7	42.7	42.7	-	0.25	0.22
14	matchbox	7.2	52.0	37.0	13.5	-	0.48	0.42
15	light bulb	28.3	103.5	60.4	60.4	9.5	0.3	0.26
16	chocolate bar	55.7	101.0	30.5	22.0	-	0.3	0.29
17	folding rule	101.4	235.0	16.4	39.0	-	0.31	0.29
18	voltage tester	16.9	139.3	17.5	15.0	11.5	0.26	0.22
19	eraser	17.9	49.0	10.3	24.2	-	1.1	0.8
20	bunch of keys	98.6	69.0	26.0	17.0	-	0.3	0.26
21	pencil	4.4	175.0	7.6	6.9	5.5	0.3	0.29

Table 4.1: Characteristics of the benchmark objects. The dimensions of three form-variable objects, the tissue pack, the paper ball, and the chocolate bar, are “dimensions of grasp”. The seventh column lists the offset of the centre of mass (CM) from the geometric centre in longitudinal direction.

These characteristics are listed in Table 4.1 and were measured or experimentally determined for all benchmark objects. The method for determining the coefficients of friction is described below. Afterwards, the number and the definition of object dimensions are discussed. The concepts of ”dimensions of sight” and ”dimensions of grasp” are introduced because form-variable objects need special treatment.

Determination of the Coefficients of Friction

The static and dynamic coefficients of friction between all objects and the desktop surface were determined experimentally by applying Newton’s laws of motion (for example, see Tipler and Mosca [2004]). The surface material is the common cotton fabric with which the desktop is covered.

In order to move an object, a force $|\boldsymbol{F_a}|$ parallel to the desktop has to be applied that is larger than the static friction $|\boldsymbol{F_s}|$ (see Figure 4.3). $|\boldsymbol{F_s}|$ is the product of the **coefficient of static friction** μ_s and the normal force $|\boldsymbol{F}_N|$ (the force exerted perpendicular to the desktop caused by the object’s mass)

$$|\boldsymbol{F}_a| \geq |\boldsymbol{F_s}| = \mu_s |\boldsymbol{F}_N| \,.$$

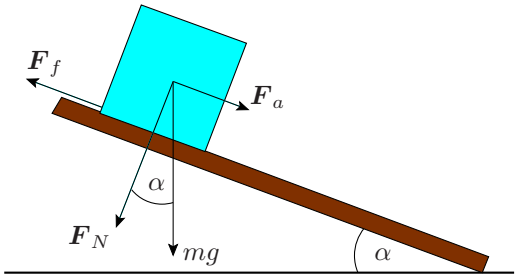

Figure 4.2: Forces acting upon an object on an inclined plane.

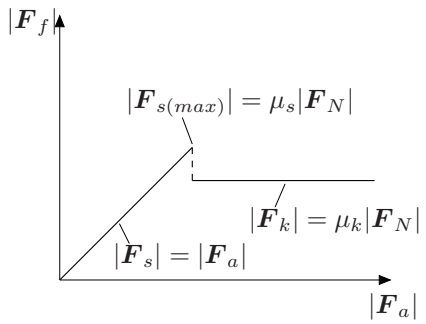

Figure 4.3: Schematic diagram of the frictional forces, as a function of the force applied on the object in parallel to the desktop.

For an object with mass m lying upon an inclined plane (see Figure 4.2; here the static friction $\boldsymbol{F_s}$ acts as the frictional force $\boldsymbol{F_f}$) this parallel force due to gravity g is

$$|\boldsymbol{F}_a| = mg \sin\alpha \,,$$

and the normal force is

$$|\boldsymbol{F}_N| = mg \cos\alpha \,.$$

Therefore,

$$mg \sin\alpha \geq \mu_s mg \cos\alpha$$
$$\tan\alpha \geq \mu_s \,.$$

That means that the coefficient of static friction is equal to the tangent of the angle at which the object begins to slide. Thus, to determine the coefficient of static friction, an inclined plane covered with the surface material of the desktop is raised in small steps until the object starts to move.

The dynamic friction $\boldsymbol{F_k}$ (now acting as frictional force $\boldsymbol{F_f}$ in Figure 4.2) resists the motion of one surface sliding over the other. It is almost constant over a wide range of low speeds (illustrated as an idealised horizontal line in Figure 4.3) and is defined by the product of the **coefficient of dynamic friction** and the normal force

$$|\boldsymbol{F}_k| = \mu_k|\boldsymbol{F}_N| \,.$$

Upon an inclined plane, the resultant parallel force $\boldsymbol{F}_p$ taking effect on a sliding object is given by

$$\boldsymbol{F}_p = \boldsymbol{F}_a + \boldsymbol{F}_k \,,$$

and its absolute value can be calculated as

$$|\boldsymbol{F}_p| = mg \sin\alpha - \mu_k|\boldsymbol{F}_N|$$
$$ma_p = mg \sin\alpha - \mu_k mg \cos\alpha \,.$$

This equation can be solved to obtain the coefficient of dynamic friction

$$\mu_k = \tan\alpha - \frac{a_p}{g \cos\alpha} \,.$$

If an object is sliding on an inclined plane with constant speed ($a_p = 0$), the coefficient of dynamic friction is equal to the tangent of the angle of the plane. Therefore, a suitable way to determine the coefficient of dynamic friction is to raise an inclined plane, covered with the material of the desktop, in steps and gently set the object on it into motion. If the object speeds up, the angle of the inclined plane has to be decreased until the object moves down the inclined plane at constant speed.

Object Dimensions

The dimensions of the benchmark objects are easily described in the case of the simple geometric shaped objects. The most simple geometric shaped objects are spherical objects, like the tennis or the golf ball because they only need one parameter, i.e. the diameter or the radius, to be fully defined in their dimensions. This shape definition is only coarse if the concavities in the surface are taken into consideration, but for grasping, this aspect can be neglected.

Two parameters are required to describe cylindrical shaped objects like, for example, the tea light or the can. For the latter, the simplification of flat head and bottom has to be accepted, which does not correspond to reality, but is sufficient for grasping because no finger contacts those regions. Three dimensions (length, width, and height) are needed to describe cuboid shaped objects like, for instance, the eraser or the matchbox.

More complex shaped objects having holes, convexities, and concavities like the adhesive tape, for example, are more difficult to describe, and a description of the form-variable paper ball is almost impossible. Therefore, and because a simplified object model in many cases is sufficient for grasp optimisation in simulation, the description of the shape of the objects is limited to a specification of three dimensions. For unification the shapes of spherical and cylindrical objects are also described by the three parameters, length, width, and height, in Table 4.1.

To determine the three dimensions of each object, the maximum distance in each direction was measured by a calliper, so that these dimensions define the smallest cuboid the respective object could be contained in. Each height is specified by the orientation the objects adopt in Figure 4.1. For instance, the heights of the can and the board marker are equal to their diameters. If the sides of an object differ in their dimensions, the larger one always defines the length and the smaller one determines the width.

Dimensions of Form-Variable Objects

For object simulation or obtaining an adequate substitute of a benchmark object, in the case of form-variable objects, it is reasonable to use ”dimensions of grasp” instead of ”dimensions of sight”. *Dimensions of sight* are dimensions that are visible and are measurable with an optical device like, for example, a 3D-scanner. They can also be measured with a calliper, but it is essential that the object is not deformed during the measurement. *Dimensions of grasp*, on the other hand, are the dimensions that a human or a robot hand adapts to, when the fingers enclose the object and the contacts are stabilised. That means that a form-variable object is deformed to the dimensions of grasp before the forces exerted on the object are strong enough to lift it.

no.	object	dimensions of sight (mm)			dimensions of grasp (mm)		
		length	width	height	length	width	height
5	tissue pack	108.0	54.0	27.0	107.0	53.0	25.0
7	paper ball	84.0	68.0	55.0	68.0	58.0	46.0
16	chocolate bar	130.0	35.0	24.0	101.0	30.5	22.0

Table 4.2: *Dimensions of sight* and *dimensions of grasp* of three form-variable benchmark objects.

The dimensions of grasp of the tissue pack and the chocolate bar are easily measured by utilising a calliper. Because the packages of these objects contain some air, the calliper jaws are forced into contact with the object until the jaws exert enough force to lift it. In the case of the paper-ball, the determination of the dimensions of grasp is more complicated because these dimensions vary with the forces the hand exerts on the object. The paper-ball can be grasped with a gentle or a powerful

grasp. To determine the dimensions that are most reasonable for the simulation of the object (see also Section 6.1.2), it was grasped by the real TUM Hand and the ”real” dimensions of grasp were measured. These are the dimensions that the paper-ball is compressed to by the hand when each of the three fingers exerts a force of about 0.7 N. The dimensions of sight are compared to the dimensions of grasp for these three objects in Table 4.2.

The remaining two form-variable objects consist of rigid components. When grasping the folding rule, its elements do not get displaced. Therefore, its dimensions of sight are equal to the dimensions listed in Table 4.1. In the case of the bunch of keys, the dimensions of grasp are those determined when the keys are pushed together. Because this alignment of the keys can be observed during the trial of grasping the bunch of keys with a robot hand, it is reasonable to use these dimensions in simulation (see Section 6.1.2). Hence, the dimensions of grasp match the dimensions of sight in the case of an aligned bunch of keys. In fact, the dimensions are similar to those of the largest key and are listed in Table 4.1.

4.1.3 Benchmark Test

We define a benchmark test for the evaluation of robot grasping. By applying this test, the grasp capabilities of different robot hands, grasp strategies and grasp types can be compared. The benchmark test is a quantitative experiment which determines the success rates when grasping objects with a real robot hand. It is applied in a scenario in which a manipulator consisting of a robot hand mounted on a robot arm has the ability to reach a target object lying on a desktop.

The following constraints and rules provide the basis for a reliable comparison of the grasp performance of different manipulators and different grasps:

Three **constraints** are associated with that framework:

- 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.

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 in Section 2.2.4, 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.

4.2 First Grasp Evaluation

The proposed benchmark system permits the evaluation of the grasp types t_1 to t_5 realised by the standard grasps g_1 to g_5 described in Section 3.2. In a grasp experiment specified by the benchmark test, we grasped every benchmark object with the standard grasps realising the basic

grasp types. The success rates of all grasp type / object pairings are used to identify the most suitable grasp type for the respective object. This facilitates the development of optimal grasps for each particular object as described in Section 6.3.

4.2.1 Choice of the Grasp Type

After developing the standard grasps g_1 to g_5 and implementing them into a robot hand setup, each of these grasps can be applied for grasping any object. Because they provide different precision and power capabilities, some of these grasps are expected to lead to more success than others when being applied for grasping a specific object.

From studies of human grasping (see, for example, Iberall and Lyons [1984], or Cutkosky and Wright [1986]), it is known that humans choose a grasp type dependent on three constraints: 1) the intended object usage (task requirements); 2) the intrinsic object properties (size, shape, weight); and 3) obstacles in the environment (workspace requirements). In knowledge-based approaches to robot grasping, features like object shapes are primarily related to grasp types depending on heuristic rules (see, for example, Stansfield [1991] or Miller et al. [2003]). In our robot hand setups, vision modules allow the observation and 3D-identification of the human hand posture. This posture is subsequently mapped to one of the implemented grasp types which can be chosen to apply with the robot hand [Steil et al., 2004]. But a different grasp type may be more suitable than the one chosen in this manner.

For choosing the grasp type most appropriate for grasping a particular object, we use the benchmark test presented in Section 4.1.3. Each available grasp type is evaluated, and the grasp type leading to the largest number of successful grasp trials is the most suitable one for grasping that specific object. If two or more grasp types lead to the same grasp success, the grasp type is chosen that is most robust against rotation around roll, pitch, and yaw axes of the hand. If in the case of two grasp types, both the number of successful grasp trials and the number of trials in which the subsequent rotations do not result in the loss of the object are equal, the grasp type is chosen which, per definition of Section 2.3.2, is more appropriate for the target object. For example, the all finger precision t_1 is the most appropriate grasp type if the matchbox (object no. 14) has to be grasped, which is a light and middle-sized object. For grasping the large and roundish tennis ball (object no. 6), the power t_3 grasp type has to be chosen. Table 4.3 provides more information and is explained in the following discussion of the first evaluation results.

4.2.2 Evaluation Results

At the time when we conducted the first grasp evaluation experiment, the Shadow Hand was not at our disposal. Thus, we only evaluated the grasp types realised for the TUM Hand. Table 4.3 presents the results of the quantitative experiment in which the 21 benchmark objects were repeatedly grasped. Grasping of each object was attempted ten times with each of the four grasp types defined in Section 2.3.2. Numeric entries show the number of successful trials for the particular grasp type / object pairing, while dashes indicate infeasible pairings.

The evaluation of the experiment reveals some interesting details. Six objects were successfully grasped in all of the ten trials when the most suitable grasp type was used. The toy propeller (object no. 2), is of a particularly complex shape and cannot be grasped by any of the four standard grasps g_1 to g_4. For this reason, the grasp type t_5 specialised for this object was realised by developing the three finger special g_5. The abduction of the index and the middle finger in pre-grasp posture are primarily responsible for all of the ten grasp trials being successful. The bunch of keys and the

no.	object	all finger precision t_1	two finger precision t_2	power t_3	two finger pinch t_4	grasp stability
1	adhesive tape	+	–	10	–	+
2	toy propeller	10 successful trials with three finger special t_5				+
3	toy cube	+	–	(+)	10	+
4	can	–	–	10	–	0
5	tissue pack	+	–	10	–	+
6	tennis ball	+	–	10	–	+
7	paper ball	(+)	–	9	–	+
8	sharpener	8	(+)	+	(+)	+
9	remote control	–	–	8	–	5
10	cup	–	–	9	–	+
11	board marker	–	7	–	–	5
12	tea light	6	–	–	–	+
13	golf ball	–	–	7	–	+
14	matchbox	7	(+)	+	(+)	+
15	light bulb	(+)	–	6	–	4
16	chocolate bar	5	–	–	–	4
17	folding rule	–	4	–	–	3
18	voltage tester	–	3	–	–	2
19	eraser	–	4	–	–	+
20	bunch of keys	0	–	–	–	0
21	pencil	–	0	–	–	0

Table 4.3: Results of the experiment for evaluating the basic grasp types t_1 to t_4 realised by the standard grasps g_1 to g_4 developed for the TUM Hand. The number of successful trials (10 to 0) out of 10 grasp attempts, denotes the most suitable grasp type. The remaining grasp types are indicated as "+": also possible, "(+)": possible but with less chances of success, "–": not possible. The toy propeller (no. 2) needs a specialised grasp type t_5. The final column gives the number of trials where the grasp is robust against rotation of the hand after lifting up ("+": robust in all trials).

pencil (object no. 20 and 21) are very close to the limits of the hardware capabilities of the TUM Hand, and no grasp trial was successful.

The standard grasps developed in accordance to the rules established in Section 3.2 in fact are suitable to grasp objects possessing characteristics as specified in Section 2.3.2. That means that for large, roundish, and heavy objects the power t_3 grasp type leads to most success, while for small or thin objects the two finger precision t_2 is most suitable. In between these two groups of objects, middle-sized objects are best grasped with the fingertips having contact, by utilising the all finger precision t_1 grasp type. Because there are no obstacles in the test scenario in which objects are grasped from a flat desktop, the kind of two finger pinch t_4 grasp type we developed in the case of the TUM Hand does not benefit from its advantages it has in constrained environments. Hence, it is the most suitable grasp type only when grasping the toy cube (object no. 3).

Altogether, the first evaluation of the standard grasps developed for the TUM Hand shows that all objects, except for two, can be grasped by utilising our grasp strategy. We observed that in most grasp trials which failed, the grasping fingers touched the object in an asynchronous manner. We discovered that as the time between the first and the last finger touching the object increases, the chance of success decreases. This qualitative assumption led to the idea of conducting an experiment on human grasping which is described in Chapter 5. Additionally, depending on the target object, it seemed that a different thumb grasp posture could lead to more success. We utilise these experiences in the optimisation strategy described in Section 6.3.

5 Experiment on Human Grasping

For optimising robot grasping, we pursued the line of biologically motivated approaches to grasp synthesis, grasp strategy, and grasp taxonomy by conducting an experiment on human grasping. The experiences gained during the performance of the preliminary robot grasp experiment described in Section 4.2.2 lead to the assumption that the developed standard grasps implemented in both robot hand setups can be improved in their chance of grasp success by optimising for touching the target object with the grasping fingers in a more simultaneous manner. To substantiate this assumption, we performed an experiment on human grasping in which we investigated the thesis that people strive for *contact simultaneity* when grasping objects, i.e. humans try to touch the target object at the same point in time with each grasping finger. Because two opposing forces can be sufficient for grasping an object (compare Section 2.2.3), we assume that people strive for contact simultaneity of two opposing fingers primarily. Both assumptions are investigated in an experiment where subjects were asked to grasp different objects.

The fact that all grasping fingers do not touch the object at the same point in time does not mean that people do not strive for contact simultaneity. Therefore, a *measure of contact simultaneity* is required. To this end, we define four different potential measures and examine which are more reliable for providing a measure of contact simultaneity. Besides analysing the assumption that people strive for contact simultaneity, a reliable measure of contact simultaneity can be used for measuring the success of a strategy that optimises the standard grasps of our robot hands for contact simultaneity.

In summary, the purposes of the experiment on human grasping are:

(i) to prove the assumption that people strive for contact simultaneity, and

(ii) to identify a reliable measure for contact simultaneity.

The potential measures and other definitions required are presented in Section 5.1. After reviewing studies on human grasping in Section 5.2, we propose an approach to determining the hand opening in Section 5.3 as an improvement to the ”grasp aperture” used commonly. In Section 5.4 the setup and the methods applied are described, before in Section 5.5 the results of the experiment are discussed. Our conclusions are summarised in Section 5.6.

5.1 Definitions

We define four different potential measures for contact simultaneity. These measure are the *grasp forming time* (GFT), the *grasp opposing time* (GOT) and, by relating them to the *grasp closure time* (GCT), the relative measures GOT% and GFT%.

The GFT is defined as the time required for *grasp forming*. This is the elapsed time between the first finger contact with the object and the first moment when it is touched by all fingers that are involved in the grasp (non-grasping fingers have to be determined and excluded from analysis). The GOT is the time required for *grasp opposing*, i.e. the time between the first finger contact and the first contact of a finger exerting an opposing force. If only two fingers are involved in the grasp, the GOT equals the GFT. If there are more grasping fingers, GOT is less than or equal

to the GFT. It is assumed that these grasp times tend to be shorter for rapidly executed grasps than for grasps conducted more deliberately. Therefore, we define two relative measures, GOT% and GFT%, and it is expected that these are more reliable measures for contact simultaneity. To determine these measures, GCT is defined to be the time between the point of peak hand opening (defined in Section 5.3) and the moment when the last grasping finger touches the object. Then:

$$GCT \geq GFT \geq GOT\,. \tag{5.1}$$

and

$$GFT\% = \frac{GFT}{GCT}\,, \qquad GOT\% = \frac{GOT}{GCT}\,. \tag{5.2}$$

For determining the points in time when the fingers touch the object, we used a measurement system that provides the opportunity to sample at a frequency of 150 Hz. This frequency is high enough that it is improbable that all finger contacts are detected in the same sample, even if the subjects strive for contact simultaneity. Therefore, a *time frame* t_f (in seconds) has to be determined in which finger contacts are considered to be simultaneous. For examining the effects of different t_f on the analysis of contact simultaneity, we define three different *contact strategies*.

The *both strategy* is applied if: “the thumb and any opposing finger touch the object simultaneously”. If a finger touches the object before the thumb, the object is dragged towards the thumb. In this case, the *finger drag strategy* is used, which means that: “any finger, except for the thumb, touches the object first”. A third strategy, the *thumb push strategy*, is applied when: “the thumb touches the object before any other finger”.

All grasp trials investigated are classified into one of the three contact strategies. The percentages of the three contact strategies in ratio to all grasp trials are examined in dependency of different time frames t_f. For this purpose, the time frame t_f is compared to the grasp opposing time (GOT). If $GOT \leq t_f$, the grasp trial is classified into the *both strategy*. If $GOT > t_f$, the strategy the grasp trial is classified into depends on the first finger touching the object as defined above.

5.2 Review of Studies on Human Grasping

There is a variety of studies on human grasping, but to our knowledge their exist no study that examines contact simultaneity in terms of the time elapsing between different finger contacts. We found a few studies investigating the grasp closure time, but in contrast to the experiment we conducted, in all of them only the thumb and the index finger were used for grasping. This simplifies determining the beginning of the grasp closure phase by identifying the ”peak grasp aperture” as measure for the maximum **hand opening**.

The ”grasp aperture” is defined as the distance between the thumb and the index finger. This kind of measure is simply calculated and compared, and only two markers, one affixed on the thumb and one on the index finger, are required. This simplicity and the lack of advanced measurement systems seem to be the reasons why the majority of existing grasping studies conducted in humans and monkeys use the peak grasp aperture as a measure for the maximum hand opening (e.g. Gentilucci et al. [1996]; Mason et al. [2004]; Paulignan et al. [1997]; Roy et al. [2002]; Saling et al. [1998]). In some human grasping experiments (e.g. Bootsma et al. [1994]; Chieffi and Gentilucci [1993]; Weir et al. [1991]; Smeets and Brenner [2001]; Zaal and Bootsma [1993]) subjects were requested to use the thumb and the index finger only. But if more than these two fingers are used, the grasp aperture as a measure for the hand opening ignores the influence of the other fingers and might be inadequate or at least inaccurate.

Mamassian [1997] found that in his experiment, where objects were grasped with all five fingers, the analysis of the time to peak grasp aperture revealed identical results if the grasp aperture was defined as the distance between the thumb and the middle finger instead of the thumb and the index finger. But although Mamassian could measure the positions of thumb, index and middle finger simultaneously, he used simply the distance between the thumb and one other finger for identifying the hand opening. Santello and Soechting [1997] not only compared the distance between thumb and index finger with that of the thumb and middle finger, but also with the distances of thumb and ring finger and thumb and little finger. The authors termed these distances ”finger span” and one result of their experiment was that the thumb-little finger combination was less accurate than the other finger combinations in the task of matching sizes of different cubes. Yokokohji et al. [2004] extended the definition of the ”finger span” by considering a third finger. The purpose of their experiment was to investigate the velocity of hand opening by using a data glove, and ”span” was defined as the area of the triangle between the three tips of the thumb, the middle finger, and the index finger.

Supuk et al. [2005] propose a method for grasp evaluation that employs all five fingers and is based on the definition of a ”pentagon plane”. By selecting three fingers, this pentagon plane is defined, while fingertip positions of the remaining two fingers are projected onto the plane perpendicularly. The area specified by this pentagon is used as a measure for hand opening, and Supuk et al. assume that this approach is more informative than the grasp aperture used commonly. But some information is still lost because the distance of the position between each of the remaining fingertips and the pentagon plane, and thus a part of their contributions to the hand opening, is neglected. An approach to determining the hand opening that also considers this information and thus is more reliable we propose in Section 5.3.

Many experiments can be found studying grasp types in monkeys (e.g. Butterworth and Itakura [1998]; Hopkins et al. [2002]; Roy et al. [2002]; Mason et al. [2004]), but less attention to this issue has been paid in studies of human grasping (e.g. Newell et al. [1993]; McDonnell et al. [2005]). The most interest in comparisons of different grasp types can be found in the area of children’s handwriting (e.g. Burton and Dancisak [2000]; Schneck and Henderson [1990]; Yakimishyn and Magill-Evans [2002]). However, these studies describe the grasp patterns used by children in holding a pencil, a crayon, or a marker and do not consider the grasp closure phase, in which the writing implement is touched. Thus, the **contact strategy** while acquiring the object was not observed, and studies which examine this issue can rarely be found.

A study that addresses both the variations in human grasp types and in the finger contact strategy was conducted by Wong and Whishaw [2004]. The authors investigated precision grasps in experiments with children and adults. The five categories of contact strategies Wong and Whishaw define can be mapped to the three categories we propose with respect to the opposing finger which touches the object first. Considering this reduction of categories, their results can be interpreted as indicating a significant strategy preference in favour of using the *finger drag* (47 %), versus the *both* (30 %), and the *thumb push* (23 %) strategy. The classification of the grasp trials was done by analysing a video of the experiment recorded by a digital camcorder, but Wong and Whishaw [2004] provide no information how precise finger contact times could be determined and in which time frame t_f two opposing contacts are considered to be simultaneous. Another result of their study was that no significant correlations were found between the sequence of finger contacts and the grasp type used. Consistent with that finding, in our experiment we do not distinguish between different grasp types, since we assume that humans strive for contact simultaneity independently of the grasp type applied.

To our knowledge there are no studies investigating **grasp times** like the grasp forming time (GFT) and the grasp opposing time (GOT), but some studies considered the grasp closure time (GCT). Topics of these investigations are hand injuries and neurological illnesses (e.g. Simoneau

et al. [1999]) and the relationship between the transport and grasp components of prehension movements (e.g. Gentilucci et al. [1992]). In the latter domain, Alberts et al. [2002] as well as Saling et al. [1998] conducted experiments in which the subjects had to avoid different obstacles during the reaching phase. Common findings were that between different obstacle conditions grasp closing distance remained constant, while grasp closing time (GCT) varied significantly. They concluded that for coordinating the components of human grasping, a spatial control seems to be more important than a temporal control.

Another study examining effects of shape and size of objects on transport and grasp times was conducted by Zaal and Bootsma [1993]. The results of their experiment reveal that GCT is not affected by the size but by the shape of the target object, in that more severe constraints on the surface area available for contact lead to longer grasp closure times. Bootsma et al. [1994] did not examine the GCT directly, but they found that movement time increases with an ”index of grasping difficulty”. This index depends on the location of the object and on its width (rectangular wooden blocks). The objects used in our experiment also show different constraints on the surface area, and the potential index of grasping difficulty is high for complex-shaped objects like, for example, the toy propeller. Therefore, we prove the GCT of all grasp trials performed by the subjects in our experiment on significant differences over all test objects (see Section 5.5), although this is not the purpose of our study.

5.3 Approach to Determining the Hand Opening

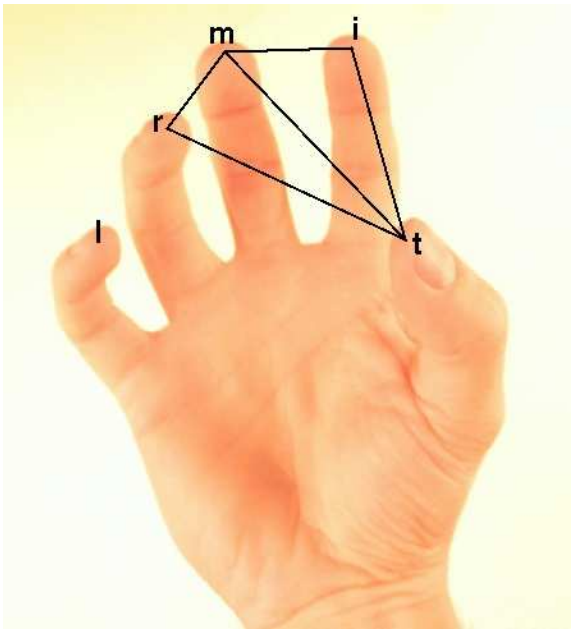

Figure 5.1: Triangles to determine the hand opening.

For determining the hand opening, we propose to use the sum of adjacent areas. These are the areas of the triangles formed by the tip of the thumb and the tips of two other adjacent fingers. In our experiment, four fingertip positions have been measured and the maximum two areas are summed. These are the areas of the triangles shown in Figure 5.1 (t–i–m and t–m–r).

One advantage of this approach is that it can be extended to the measurement of five fingers by

simply adding the area of one more triangle (t–r–l). Compared to the ”pentagon plane” approach of Supuk et al. [2005], our approach is more reliable. By summing up areas of different planes, the three-dimensional vectors of the fingertip positions are considered, and no information is lost (compare Section 5.2).

The approach we propose requires that the fingers the subjects use in each grasp trial have to be noted. Then the *hand opening* is determined in dependency of the grasping fingers:

- two fingers: *distance* between the thumb and the second finger (distance between t–i, t–m, or t–r),
- three fingers: *area* of the triangle between the thumb and the two fingers remaining (area of t–i–m or t–m–r),
- four fingers: *sum of areas* of the two triangles between the thumb and two adjacent fingers (areas of t–i–m and t–m–r).

The *peak hand opening* is the maximum hand opening that can be measured during a grasp trial. The point in time when this peak hand opening occurs determines the beginning of grasp closure and has to be identified for calculating the grasp closure time (GCT).

5.4 Setup and Methods

Subjects

A total of twenty-one adult volunteers (eleven females, ten males) participated in the experiment, their age ranging from 24 to 50 years (mean age 31 years). Although not preconditioned, all subjects were right-handed. Thus, a comparison of the experimental results between right- and left-handed participants was not feasible. Nevertheless, differences in these results are not expected. The volunteers were naive as to the purpose of the experiment until they had finished their trials.

Apparatus

For the apparatus, four everyday objects (board marker, tea light, toy propeller, can) were chosen from the benchmark system for grasp evaluation proposed in Section 4.1, aiming for large differences in size and shape of the objects. Additionally, each of these objects is grasped with a different grasp type when grasped with our robot hands. Thus, a representative set of objects is provided that is suitable for the task of investigating grasping of everyday objects. The objects were placed on a flat table by aligning them in a semicircle onto marked positions with the fifth marked position (initially the leftmost from subject’s view) unoccupied. Each object was equipped with one infrared LED-marker.

The subjects were seated (comfortably) on a chair in front of the table (width = 78 cm, length = 78 cm , height = 75 cm) (see Figures 5.2a and 5.2b). The dominant hand of each subject (the right hand) was covered with a latex glove prepared with four additional LED-markers. The main reason for leaving one fingertip unobserved was that each chain of LED-markers only consisted of four LEDs. Furthermore, preliminary experiments have shown that occlusions of a marker on the little finger were responsible for most grasp trials that were unusable for data analysis. Other studies on the analysis of human grasping and manipulation skills revealed that the observation of

(a) Overview of the Apparatus. (b) LED-markers on glove and objects.

Figure 5.2: Apparatus.

the little finger is not necessary even when investigating finger manipulation patterns (for example, see Kurita et al. [2004]). Therefore, the little finger was not covered with an LED-marker.

A valuable property of the latex material is that it is formfitting to different sizes of hands and fingers. Therefore, the LED-markers on the glove could easily be positioned tightly over the covered fingernails. For hygienic reasons and to reuse this prepared glove for all participants, another thin latex glove was pulled over the hand first. In addition, this affected the quality of fit of the LED-markers positively. Analysing data of a latex glove prepared in this way has advantages over the utilisation of a common data glove. Because the fingers and objects are equipped with the same kind of LED-markers, their positions and velocities can be measured with the same measurement system and thus can directly be compared. With a data glove, at least three additional markers attached onto it would be required to determine its position and orientation providing the possibility to compare the data of fingers and objects. However, the accuracy of the obtained positions and velocities, converted from one measurement system into another, is questionable.

Recording System

Movements of fingers and objects were recorded using the 3D motion tracking system AS 200 by LUKOtronic (Lutz-Kovacs-Electronics OEG, Innsbruck, Austria). The main component of this tracking device is a portable measurement system consisting of three infrared-cameras fixated on a hinged mount. In the opened state (see Figure 5.3c), the relative position and orientation of the cameras are defined exactly. Each camera measures the position of the LED-markers which alternately emit light in the infrared band of radiation. The three-dimensional coordinates of the infrared LED-markers are calculated from the three two-dimensional camera-coordinates with a resolution of 0.1 mm.

(a) Compact. (b) Half-open. (c) Open.

Figure 5.3: Hinged measurement system consisting of three cameras.

Preliminary tests with this measurement system showed that occlusions of the infrared LED-markers of different fingers led to the rejection of many grasp trials. Fortunately, the tracking system AS 200 provides the possibility to interconnect a second measurement system with an additional three infrared-cameras, which was positioned at a distance of 2.6 m to the first system and with a change of 105 degrees in orientation. (Figure 5.2a presents an overview only and does not show the correct positions and orientations). The height of both measurement systems was 1.21 m, and the tilt angle was 16 degrees. With this setup, a total of eight infrared-LED-markers were measured at a sampling frequency of 150 Hz each, and in most grasp trials no occlusions occurred.

Procedure

To ensure that all subjects were provided with the same preliminary information, each participant was instructed by a one page handout describing the task of the experiment. The experiment comprised a pick-and-place task. Subjects were asked to take an object next to the unoccupied position marker and place the object upon it. With the task of placing the object onto a different position, the attention of the participant is not focused on the task of grasping it solely. We assume that the result is more unconscious and therefore a more natural way to grasp the object. For the same reason, the goal of the experiment was concealed before all trials were completed. The only constraint was avoiding a rotation of the objects during the pick-and-place task. This was required for determining accurate times of finger contacts as described below.

The first object to be grasped was the board marker, which had to be laid down onto the marked position to the left (from subject's view). The next object (tea light) had to be placed onto the prior position of the board marker. Then the propeller and afterwards the can had to be grasped and placed one position further to the left. Now the last object moved (can) had to be grasped again and laid down onto its prior position right aside.[1] The remaining objects had to follow in the reverse order, each one marked position to the right, until the leftmost position was unoccupied again. This finished the first cycle of the experiment, and the participant had to start all over. Each subject had to complete five of these cycles.

After reading the instructions, the subjects had the possibility to ask questions about the task. Only a few subjects took this chance to ask how fast the pick-and-place task had to be performed, or whether there was a special way in which the object should be grasped. Both questions were answered by advising these subjects to carry out the task in a way as natural as possible.

Determining the Finger Contacts

To determine the points in time when the fingers touch the target object, the velocities of fingers and objects are used, which have to be calculated from the positional data. The high number of four fingers observed is the reason that a more advanced strategy has to be found for detecting finger contacts compared to almost every grasping experiment existing. Most of these experiments investigate two finger grasping and detect the lifting of an object by using a binary contact sensor. Others assume that the fingers have contact when the object begins to move (for example, see Zaal and Bootsma [1993]). Smeets and Brenner [2001] use a velocity threshold for indicating the end of the finger movements before touching the target object. But a target object might be in motion before a finger has contact due to other fingers touching the object before.

[1] The procedure in this state is shown in Figure 5.2b.

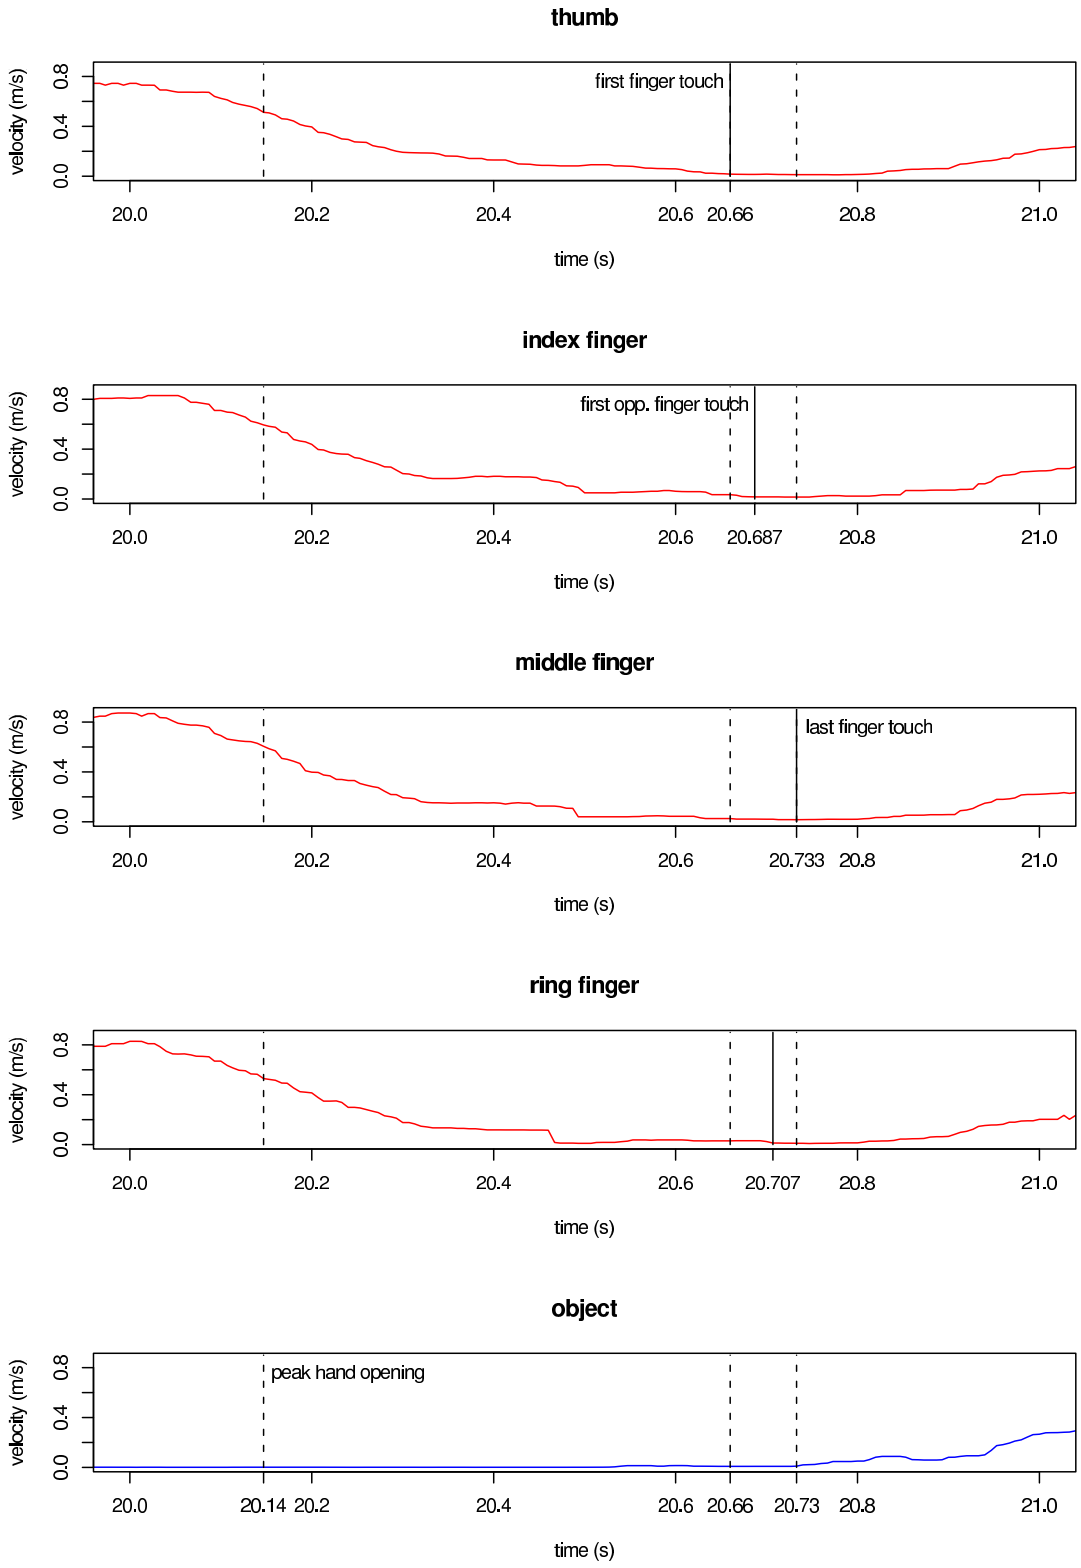

Figure 5.4: Typical example of finger and target object velocities as a function of time. Solid lines determine the contact times of the respective fingers. In this grasp trial, the first finger touching the object is the thumb, and the first opposing force is provided by the index finger. The middle finger touches the object at last. $GCT = 0.59\ s; GFT = 0.07\ s; GOT = 0.027\ s$.

We propose the following procedure for finding times of finger contact:

1. For each grasp trial find the time of the maximum object velocity indicating the end of object acceleration in lift-off phase. Then search in the direction of negative time and find the first point when all finger velocities are larger than the object velocity. The resulting two points in time determine the search range for finger contact times.

2. Search in direction of positive time and find the first points in time when the velocities of the grasping fingers are equal to the object velocity. These are the points in time when the fingers touch the object.

The main idea behind our approach to determining the finger contacts is that a contact is established if during grasp closure the velocity of the grasping finger equals the velocity of the target object. Before the object is touched, its velocity is equal to zero, while the velocities of the fingers approaching the object are larger. During the grasp trial and before actually touching the object, the velocity of a finger might become zero leading to the determination of an incorrect contact time. This potential error is avoided by the first step of the procedure proposed, in which the search range for contacts is limited.

Another potential source of error is the rotation of the object when it is grasped by the subject. If the LED-marker on a finger that touches the object has a different distance to the rotational axis than the object's LED-marker, different marker-velocities are detected. In this case, a contact is not identified although the finger touches the object already. Therefore, it was required that the participants did not rotate the objects during the grasp trials.

With this procedure for finding times of finger contact, not only finger contacts are detected before the object is moved (velocities equal zero), but also contacts that occur afterwards (velocities non-zero) can be identified. A typical example of the chronological sequences of finger and target object velocities with contact times marked is shown in Figure 5.4.

Data Analysis

During each grasp trial of the experiment, it was observed which fingers did not have contact with the target object. Afterwards, it was checked whether the LED-markers of the grasping fingers had been visible during grasp closure. Trials that showed interrupted data in the grasp closure phase were rejected, and 504 trials remained for data analysis.

High frequence noise of the three-dimensional position signals was removed by using a median filter. Calculation of the hand opening was done by using Heron's formula [Dunham, 1990] after determining the distances between the position vectors of the fingertips. The point in time of peak hand opening was taken as the beginning of grasp closure. The velocities in each direction were obtained by differentiating the position data of the x-, y,- and z-axes. The velocity of the movement along the path of the trajectory was calculated by squaring and summing up the three axial velocity values and taking the square root of the result.

The kinematic trajectories were used to derive the following dependent variables: times of finger contact, grasp closure time (GCT), grasp forming time (GFT), grasp opposing time (GOT), ratio between grasp forming time and grasp closure time (GFT%), ratio between grasp opposing time and grasp closure time (GOT%).

Analysis of variables derived from the kinematic trajectories was performed using the R-statistic tool. T-tests and ANalysis Of Variance (ANOVA) were used to compare mean values between

conditions and paired t-tests to compare mean values within conditions. Pearson correlation coefficients were calculated for selected parameters of variables. For all statistical analyses, a significance level of $p < 0.05$ was used, and a level of $p < 0.001$ denoted strong significance. Unless otherwise noted, each statistical test implies all subjects (row "*all subjects:*" in the tables of Section 5.5).

5.5 Results and Discussion

We analysed each of the four measures defined in Section 5.1:

- grasp forming time (GFT) (Time between the points when the first and the last grasping finger touches the target object.),
- grasp opposing time (GOT) (Time between the points when the first finger and the first opposing finger touches the target object.),
- grasp forming time in percentage of grasp closure time (GFT%),
- grasp opposing time in percentage of grasp closure time (GOT%),

in terms of two different characteristics: (i) if it proves or disproves the assumption that humans strive for contact simultaneity when grasping objects; and (ii) if it is a reliable measure for contact simultaneity.

Tables 5.1 and 5.2 show mean and standard deviation of grasp closure time (GCT), grasp forming time (GFT), and grasp opposing time (GOT) of all subjects grasping the four objects.

In this section, the results of the analysis of the grasp times GFT and GOT are presented at first. The GOT of a grasp trial determines whether or not a "both" grasp strategy is applied. To this end, GOT is compared with the time frame t_f. We examine the effect of different values for t_f on the ratios of the grasp strategies. Then the grasp closure time (GCT) is analysed and the results are compared to other studies investigating GCT, which are reviewed in Section 5.2.

In the experiment, it was recognised that subjects took different amounts of time to perform the grasp trials. Since we assume that the time required for grasp closure (GCT) affects GFT and GOT (see Section 5.1), we performed a correlation analysis between GFT and GCT as well as between GOT and GCT. The results are presented, before the analysis of all measures on gender dependency is discussed. Since we assume that the simultaneity of finger contacts does not depend on gender, measures that are independent of gender are more suitable measures for contact simultaneity. An analysis of GFT% and GOT% provide the concluding results.

Analysis of GFT and GOT

Analysis of variance (ANOVA) showed no significant differences in GFT ($F(3, 57) = 1.359$) and GOT ($F(3, 57) = 0.41$) between all objects ($p > 0.05$ in each case). Mean and standard deviation over all objects are 0.06 ± 0.06 seconds in the case of GFT and 0.03 ± 0.04 seconds in the case of GOT (compare Table 5.4). These values are very small. We conclude that both GFT and GOT prove that the subjects strove for contact simultaneity when closing the grasp.

But since there does not exist a critical value for appointing a time frame in which contacts have to be made to be considered as simultaneous, different critical values in the form of time frames t_f after first finger contact were tested and compared to the results found by Wong and Whishaw

subject		board marker			tea light		
no.	age	GCT (s)	GFT (s)	GOT (s)	GCT (s)	GFT (s)	GOT (s)
females							
1	27	0.83 ±0.10	0.06 ±0.03	0.03 ±0.04	0.98 ±0.66	0.08 ±0.04	0.03 ±0.03
2	30	0.94 ±0.31	0.05 ±0.03	0.04 ±0.03	1.10 ±0.40	0.06 ±0.07	0.02 ±0.03
3	31	0.41 ±0.00	0.03 ±0.00	0.01 ±0.00	0.78 ±0.58	0.04 ±0.02	0.03 ±0.03
4	26	– ± –	– ± –	0.00 ±0.00	0.71 ±0.09	0.06 ±0.06	0.02 ±0.02
5	27	0.96 ±0.54	0.03 ±0.02	0.01 ±0.02	0.69 ±0.26	0.03 ±0.03	0.03 ±0.03
6	43	0.85 ±0.37	0.03 ±0.03	0.02 ±0.02	0.70 ±0.25	0.06 ±0.02	0.03 ±0.02
7	50	0.54 ±0.25	0.04 ±0.04	0.01 ±0.01	0.49 ±0.30	0.05 ±0.04	0.03 ±0.04
8	24	0.63 ±0.20	0.06 ±0.06	0.02 ±0.02	0.76 ±0.57	0.06 ±0.05	0.04 ±0.03
9	29	0.95 ±0.40	0.07 ±0.07	0.07 ±0.07	0.98 ±0.41	0.06 ±0.05	0.03 ±0.04
10	26	0.64 ±0.17	0.13 ±0.09	0.02 ±0.02	1.06 ±0.60	0.09 ±0.12	0.04 ±0.04
11	27	0.69 ±0.14	0.05 ±0.04	0.02 ±0.03	0.73 ±0.33	0.05 ±0.06	0.03 ±0.06
all females		0.76 ±0.30	0.06 ±0.06	0.03 ±0.03	0.84 ±0.45	0.06 ±0.05	0.03 ±0.03
males							
12	37	0.59 ±0.11	0.06 ±0.03	0.05 ±0.03	0.52 ±0.18	0.02 ±0.01	0.00 ±0.01
13	28	1.10 ±0.65	0.08 ±0.07	0.02 ±0.02	1.11 ±0.13	0.04 ±0.05	0.03 ±0.06
14	35	1.04 ±0.42	0.04 ±0.06	0.04 ±0.06	1.28 ±0.48	0.04 ±0.06	0.04 ±0.06
15	28	0.98 ±0.52	0.06 ±0.04	0.03 ±0.04	0.79 ±0.51	0.06 ±0.07	0.05 ±0.08
16	26	1.44 ±0.49	0.10 ±0.06	0.08 ±0.05	0.93 ±0.21	0.12 ±0.15	0.00 ±0.00
17	30	0.92 ±0.33	0.04 ±0.04	0.03 ±0.03	1.01 ±0.64	0.04 ±0.04	0.03 ±0.04
18	31	0.86 ±0.19	0.07 ±0.08	0.05 ±0.05	0.89 ±0.48	0.06 ±0.06	0.04 ±0.06
19	37	0.46 ±0.16	0.02 ±0.02	0.01 ±0.01	0.59 ±0.14	0.05 ±0.04	0.03 ±0.02
20	30	0.90 ±0.40	0.05 ±0.02	0.04 ±0.03	0.63 ±0.54	0.03 ±0.02	0.01 ±0.01
21	24	0.83 ±0.31	0.11 ±0.10	0.06 ± 0.07	1.02 ±0.35	0.13 ±0.07	0.05 ±0.04
all males		0.89 ±0.41	0.06 ±0.06	0.04 ±0.04	0.87 ±0.45	0.06 ±0.07	0.03 ±0.04
all subjects		0.82 ±0.38	0.06 ±0.06	0.03 ±0.04	0.85 ±0.45	0.06 ±0.06	0.03 ±0.04

Table 5.1: Mean and standard deviation of grasp closure time (GCT), grasp forming time (GFT), and grasp opposing time (GOT) of all subjects grasping the board marker and the tea light. Values are given in seconds (s).

[2004] (as explained in Section 5.2). When using a critical value of $t_f = 7\ ms$, the same ratio for the *both* contact strategy was obtained (30 %) (see Figure 5.5). Another analogy is the fact that the *finger drag* strategy was favoured over the *thumb push* strategy. A critical value of $t_f = 0.02\ s$ is already sufficient for obtaining a ratio of more than 50 percent for the *both* strategy, and a time frame of $t_f = 0.1\ s$ (which is still short compared to the GCT, as analysed below) involves 93 percent.

If the first grasp trials of all subjects were neglected, mean and standard deviations of GFT and GOT would be even less, and thus the critical value for the time frame t_f preserving the same percentages for the *both* strategy would be smaller. Paired t-tests between the measurements of the first and the last cycle of the grasp procedure show that there is a training effect towards smaller values ($p < 0.001$ in the case of GFT, $p = 0.028$ in the case of GOT). We conclude that this training effect also indicates that the subjects strove for contact simultaneity.

Analysis of GCT

As listed in Table 5.4, mean and standard deviation of GCT over all subjects and all objects are 0.82 ± 0.52. In opposition to GFT and GOT, paired t-tests between the first and the last measurements per object of each person showed no significant training effect for GCT ($p = 0.2566$).

subject		toy propeller			can		
no.	age	GCT (s)	GFT (s)	GOT (s)	GCT (s)	GFT (s)	GOT (s)
females							
1	27	0.66 ±0.26	0.07 ±0.06	0.06 ±0.05	0.57 ±0.53	0.03 ±0.01	0.01 ±0.02
2	30	0.43 ±0.14	0.03 ±0.04	0.02 ±0.04	0.68 ±0.30	0.08 ±0.03	0.01 ±0.02
3	31	0.49 ±0.14	0.09 ±0.08	0.03 ±0.03	0.90 ±0.71	0.08 ±0.04	0.03 ±0.02
4	26	0.72 ±0.25	0.04 ±0.04	0.02 ±0.03	1.13 ±0.88	0.03 ±0.03	0.01 ±0.03
5	27	0.58 ±0.41	0.08 ±0.05	0.05 ±0.04	0.85 ±0.51	0.04 ±0.03	0.01 ±0.01
6	43	0.38 ±0.20	0.02 ±0.01	0.01 ±0.01	0.84 ±0.33	0.04 ±0.05	0.02 ±0.01
7	50	0.47 ±0.26	0.08 ±0.12	0.01 ±0.01	0.96 ±0.28	0.04 ±0.03	0.02 ±0.02
8	24	0.48 ±0.02	0.07 ±0.02	0.00 ±0.00	0.64 ±0.54	0.06 ±0.03	0.02 ±0.02
9	29	0.21 ±0.09	0.03 ±0.05	0.02 ±0.01	1.09 ±0.41	0.09 ±0.07	0.02 ±0.04
10	26	0.59 ±0.18	0.08 ±0.07	0.03 ±0.02	0.45 ±0.11	0.05 ±0.07	0.02 ±0.01
11	27	0.37 ±0.26	0.07 ±0.09	0.05 ±0.07	0.51 ±0.13	0.05 ±0.05	0.02 ±0.02
all females		0.49 ±0.26	0.06 ±0.07	0.03 ±0.04	0.80 ±0.52	0.05 ±0.04	0.02 ±0.02
males							
12	37	0.88 ±0.13	0.07 ±0.07	0.05 ±0.05	0.60 ±0.16	0.03 ±0.02	0.01 ±0.02
13	28	0.32 ±0.40	0.00 ±0.00	0.00 ±0.00	1.45 ±1.38	0.06 ±0.04	0.03 ±0.04
14	35	1.01 ±0.50	0.14 ±0.11	0.06 ±0.06	1.76 ±0.70	0.17 ±0.09	0.06 ±0.06
15	28	0.90 ±0.51	0.10 ±0.08	0.05 ±0.04	0.60 ±0.35	0.06 ±0.05	0.02 ±0.03
16	26	0.81 ±0.24	0.20 ±0.20	0.02 ±0.01	1.32 ±1.31	0.08 ±0.08	0.04 ±0.06
17	30	0.67 ±0.49	0.04 ±0.04	0.03 ±0.03	1.07 ±0.98	0.16 ±0.09	0.06 ±0.06
18	31	0.59 ±0.31	0.07 ±0.05	0.04 ±0.03	1.16 ±0.59	0.09 ±0.06	0.04 ±0.05
19	37	0.56 ±0.26	0.08 ±0.08	0.05 ±0.04	0.91 ±0.51	0.05 ±0.03	0.03 ±0.02
20	30	1.08 ±0.26	0.08 ±0.09	0.05 ±0.04	1.26 ±1.42	0.03 ±0.01	0.03 ±0.01
21	24	0.75 ±0.38	0.08 ±0.14	0.03 ±0.04	1.45 ±0.42	0.14 ±0.08	0.06 ±0.06
all males		0.76 ±0.39	0.08 ±0.08	0.04 ±0.04	1.11 ±0.86	0.08 ±0.07	0.04 ±0.04
all subjects		0.63 ±0.36	0.07 ±0.08	0.04 ±0.04	0.96 ±0.72	0.07 ±0.06	0.03 ±0.04

Table 5.2: Mean and standard deviation of grasp closure time (GCT), grasp forming time (GFT), and grasp opposing time (GOT) of all subjects grasping the toy propeller and the can. Values are given in seconds (s).

ANOVA over all objects showed strong significant differences in GCT ($F(3, 57) = 8,285$, $p < 0.001$). Two-sample t-tests of the toy propeller versus each of the other three objects revealed that less time was required ($p < 0.001$ in each case), and GCT of grasping the can was significantly larger than that of the other objects ($p < 0.001$ in each case).

In the case of the can, we assume that the reason why more time was needed for grasp closure is that the distances between the fingertip positions in pre-grasp posture and in grip posture are larger than when grasping the other objects. These distances can be very small when using a precision grasp, in which only the fingertips touch the object. But most subjects preferred to apply a power grasp for grasping the can. To encompass the object, this grasp type requires a large peak hand opening. That means that the fingertips in pre-grasp posture are far from their position in the grip posture, in which the fingers enclose the can. Additionally, in the reaching phase, a larger distance has to be covered in the approach direction for providing a contact between the palm and the target object, after the peak hand opening has to be reached. Compared to a precision grasp, this movement requires more time too.

In contrast to Zaal and Bootsma [1993] and Bootsma et al. [1994], the GCT was shorter in the case of one object with a higher "index of grasping difficulty" (see Section 5.2). In both studies referred to, only small or simple-shaped objects were used. But the complex shape of the toy propeller together with its large size is responsible for that subjects adapting their pre-grasp posture at the end of the reaching phase. The decision of "what is the best grasp" and thus the time of peak hand

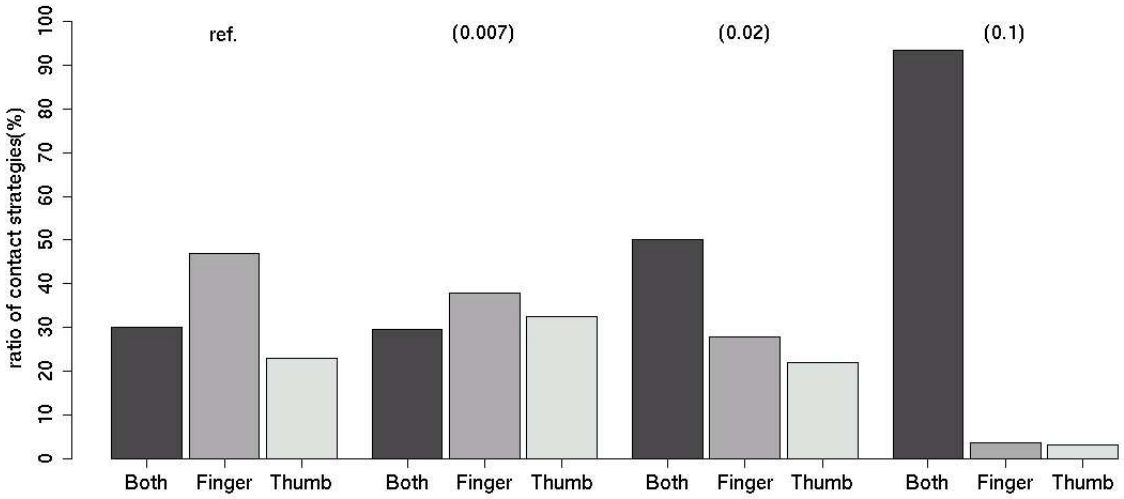

Figure 5.5: Ratios of simultaneous grasps in percentages of all grasp trials represented by the *both* contact strategy and compared to the *finger drag* and the *thumb push* strategies. The "level of simultaneity" depends on the critical value of the time frame t_f (0.007 s, 0.02 s, 0.1 s) (ref.: Wong and Whishaw [2004]).

opening occurs much later, and this shortens the GCT.

subj.	board marker		tea light		toy propeller		can	
no.	GFT%	GOT%	GFT%	GOT%	GFT%	GOT%	GFT%	GOT%
females								
1	7.6 ±3.5	3.3 ±4.1	10.2 ±6.7	4.9 ±5.3	9.2 ±5.8	8.2 ±4.7	7.9 ±4.1	0.6 ±1.4
2	6.7 ±5.6	3.9 ±2.3	6.9 ±7.8	2.2 ±2.9	6.7 ±6.4	4.0 ±6.5	13.8 ±6.9	1.3 ±1.3
3	6.4 ±0	3.2 ±0	5.8 ±3.8	3.3 ±2.3	16.9 ±10.0	5.4 ±3.4	12.1 ±9.5	3.4 ±3.9
4	– ± –	– ± –	8.1 ±8.0	3.6 ±2.4	5.8 ±4.2	2.9 ±2.5	3.4 ±3.8	0.9 ±1.0
5	3.6 ±3.8	1.5 ±1.8	4.4 ±3.6	4.4 ±3.6	17.7 ±12.8	8.8 ±5.2	5.1 ±3.1	1.6 ±1.5
6	4.4 ±3.7	3.5 ±3.6	8.3 ±2.6	4.1 ±3.4	6.2 ±3.4	3.9 ±2.6	5.1 ±3.7	2.7 ±1.1
7	8.9 ±7.8	3.3 ±3.4	9.4 ±6.9	4.6 ±4.4	17.3 ±23.0	2.2 ±0.7	4.2 ±2.6	1.7 ±1.7
8	9.6 ±7.6	3.8 ±4.8	9.3 ±7.4	5.4 ±4.3	15.3 ±4.4	0.6 ±0.9	14.8 ±8.5	6.2 ±6.5
9	7.2 ±5.4	6.4 ±5.9	7.0 ±7.0	3.0 ±2.8	12.4 ±11.1	7.4 ±1.8	7.9 ±5.3	2.0 ±2.7
10	21.1 ±14.9	3.1 ±2.7	11.4 ±18.0	5.1 ±6.1	11.8 ±9.0	4.8 ±4.9	10.4 ±8.7	5.1 ±3.0
11	7.6 ±6.7	3.2 ±3.2	5.7 ±4.3	3.6 ±5.0	33.7 ±41.7	19.6 ±27.8	10.4 ±9.2	3.6 ±2.4
all	8.9 ±8.7	3.5 ±3.7	7.7 ±6.8	3.9 ±3.8	14.8 ±19.6	7.1 ±11.8	8.2 ±7.0	2.6 ±3.0
males								
12	9.4 ±4.7	8.6 ±5.2	4.2 ±4.3	0.9 ±1.8	8.0 ±6.5	5.2 ±4.9	4.9 ±4.2	1.9 ±2.9
13	9.8 ±8.5	2.5 ±2.9	3.6 ±4.5	2.8 ±4.9	1.7 ±2.9	1.7 ±2.9	6.6 ±6.1	4.3 ±6.6
14	3.3 ±3.6	3.3 ±3.6	4.0 ±6.1	4.0 ±6.1	11.6 ±7.4	6.1 ±4.7	11.5 ±7.4	3.8 ±3.9
15	8.3 ±10.7	2.6 ±2.6	6.0 ±5.7	4.4 ±6.0	13.1 ±9.7	6.2 ±3.4	9.3 ±6.6	3.4 ±3.0
16	8.2 ±6.6	6.9 ±5.9	15.5 ±21.8	0.5 ±0.7	21.9 ±17.8	2.4 ±2.4	8.8 ±7.9	2.6 ±1.8
17	5.5 ±5.3	4.4 ±4.4	4.1 ±3.6	3.9 ±3.3	4.3 ±3.8	3.7 ±3.8	21.6 ±9.5	7.8 ±7.8
18	9.1 ±10.4	5.7 ±5.1	5.9 ±4.2	3.4 ±3.7	16.4 ±13.1	6.7 ±5.5	10.9 ±8.1	4.2 ±4.4
19	4.7 ±2.5	2.5 ±2.2	9.3 ±6.4	5.6 ±2.7	13.5 ±10.1	7.6 ±4.4	6.3 ±4.8	3.4 ±2.6
20	6.7 ±2.8	4.6 ±3.2	6.1 ±6.3	3.0 ±4.7	7.1 ±5.8	4.0 ±3.7	8.3 ±10.6	8.3 ±10.6
21	11.5 ±9.6	6.3 ±5.0	14.9 ±9.6	5.0 ±3.9	15.3 ±23.0	6.4 ±7.1	9.7 ±4.6	5.3 ±5.7
all	7.4 ±6.8	4.7 ±4.3	7.2 ±8.1	3.5 ±3.9	10.8 ±10.7	5.3 ±4.6	9.5 ±7.7	4.0 ±4.5
all s.	8.0 ±7.7	4.2 ±4.1	7.5 ±7.5	3.7 ±3.9	12.8 ±15.9	6.2 ±8.9	8.9 ±7.3	3.3 ±3.9

Table 5.3: Mean and standard deviation of grasp forming time (GFT) and grasp opposing time (GOT) in percentage (%) of grasp closure time (GCT).

subject no.	all objects GCT	GFT	GOT	GFT%	GOT%
females					
1	0.78 ±0.45	0.06 ±0.04	0.03 ±0.04	8.82 ±5.09	4.44 ±4.82
2	0.82 ±0.41	0.06 ±0.05	0.02 ±0.03	8.08 ±7.09	2.77 ±3.90
3	0.71 ±0.53	0.07 ±0.05	0.03 ±0.02	11.30 ±8.88	3.99 ±3.13
4	0.91 ±0.64	0.04 ±0.04	0.02 ±0.03	5.04 ±4.75	2.06 ±2.13
5	0.77 ±0.43	0.04 ±0.04	0.03 ±0.03	7.73 ±8.80	4.08 ±4.34
6	0.74 ±0.32	0.04 ±0.03	0.02 ±0.02	6.14 ±3.56	3.51 ±2.81
7	0.64 ±0.33	0.05 ±0.07	0.02 ±0.02	10.11 ±13.89	2.65 ±2.46
8	0.66 ±0.41	0.06 ±0.05	0.02 ±0.02	11.00 ±7.57	4.53 ±4.83
9	0.85 ±0.48	0.07 ±0.06	0.03 ±0.04	8.44 ±7.13	4.34 ±3.98
10	0.65 ±0.34	0.09 ±0.09	0.03 ±0.03	14.35 ±13.04	4.40 ±3.72
11	0.57 ±0.27	0.06 ±0.06	0.03 ±0.05	14.67 ±24.28	7.75 ±15.77
all females	0.73 ±0.43	0.06 ±0.05	0.03 ±0.03	9.74 ±11.71	4.21 ±6.59
males					
12	0.65 ±0.20	0.04 ±0.04	0.03 ±0.04	6.47 ±5.16	3.93 ±4.71
13	1.09 ±0.89	0.05 ±0.05	0.02 ±0.04	5.88 ±6.38	3.00 ±4.60
14	1.27 ±0.59	0.09 ±0.10	0.05 ±0.06	7.42 ±7.09	4.20 ±4.42
15	0.77 ±0.44	0.07 ±0.06	0.03 ±0.04	9.22 ±7.73	4.01 ±3.66
16	1.21 ±0.93	0.11 ±0.10	0.04 ±0.05	11.44 ±12.38	3.07 ±3.54
17	0.91 ±0.60	0.06 ±0.07	0.04 ±0.04	7.83 ±8.58	4.72 ±4.82
18	0.86 ±0.44	0.08 ±0.06	0.04 ±0.05	10.62 ±9.98	5.02 ±4.71
19	0.66 ±0.36	0.05 ±0.05	0.03 ±0.03	8.30 ±6.87	4.70 ±3.46
20	0.93 ±0.51	0.05 ±0.06	0.03 ±0.03	6.84 ±5.05	4.41 ±4.34
21	0.99 ±0.41	0.12 ±0.09	0.05 ±0.05	12.96 ±11.68	5.74 ±4.89
all males	0.91 ±0.58	0.07 ±0.07	0.04 ±0.04	8.65 ±8.45	4.37 ±4.36
all subjects	0.82 ±0.52	0.06 ±0.06	0.03 ±0.04	9.19 ±10.20	4.29 ±5.57

Table 5.4: Mean and standard deviation of grasp closure time (GCT), grasp forming time (GFT), grasp opposing time (GOT), and the relative the relative measures, GFT% and GOT%, in percentage of the grasp closing time. Values for GCT, GFT, and GOT are given in seconds (s).

Correlation between GFT, GOT, and GCT

Pearson correlation analysis between GFT and GCT led to the result that there was a significant positive correlation between GFT and GCT in different tested groups: all men over all objects ($r = 0.32, p < 0.001$), all women over all objects ($r = 0.13, p = 0.037$), all subjects over all objects ($r = 0.26, p < 0.001$). A strong significant positive correlation was found between GOT and GCT in different tested groups: all men over all objects ($r = 0.36, p < 0.001$), all women over all objects ($r = 0.24, p < 0.001$), all subjects over all objects ($r = 0.34, p < 0.001$).

These results sustain the assumption that GCT affects GFT and GOT. Thus, the means and standard deviations of GFT in percentage of GCT and GOT in percentage of GCT (GFT% and GOT% are listed in Table 5.3) are more reliable measures for contact simultaneity than GFT or GOT.

Further analysis revealed a strong positive correlation between GFT and GOT in different tested groups: all men over all objects ($r = 0.63, p < 0.001$), all women over all objects ($r = 0.52, p < 0.001$), all subjects over all objects ($r = 0.6, p < 0.001$). We conclude that GFT% as well as GOT% serve as reliable measures for contact simultaneity.

Analysis of Gender Dependency

When comparing female and male subjects, two-sample t-test revealed that women required significantly less time for grasp forming (GFT) ($p = 0.028$) and grasp opposing (GOT) ($p = 0.0014$), and strongly significantly less time for grasp closure (GCT) ($p < 0.001$). A comparison of GFT% and GOT% between women and men showed no significant differences. This supports the assumption that GFT% and GOT% are more reliable measures for contact simultaneity than GFT and GOT, since it was assumed that the simultaneity of finger contacts does not depend on gender.

Analysis of GFT% and GOT%

Mean and standard deviation over all subjects and all objects are 9.19 ± 10.2 percent for GFT%, and 4.29 ± 5.57 percent for GOT% (see Table 5.4). That means that only small ratios (less than ten percent) of the time needed for grasp closure are used for grasp forming (GFT) and grasp

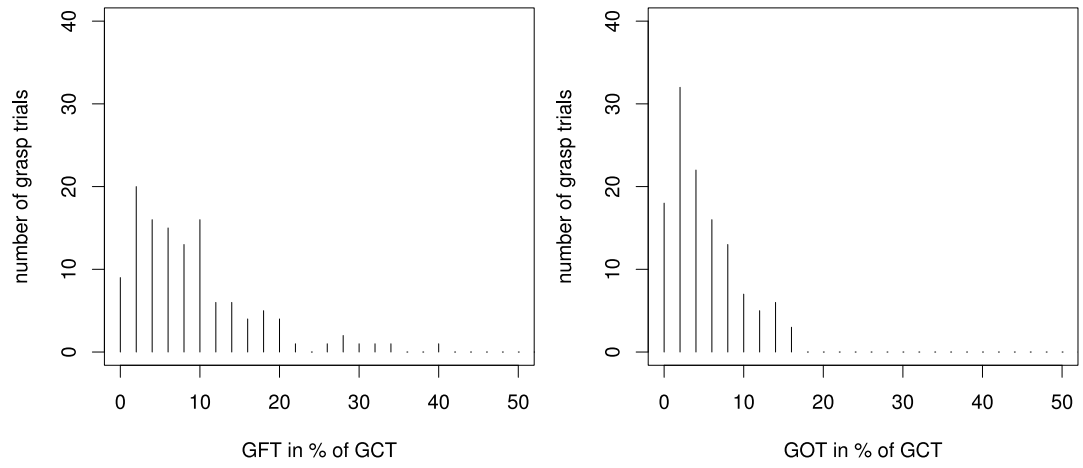

(a) GFT%; board marker.

(b) GOT%; board marker.

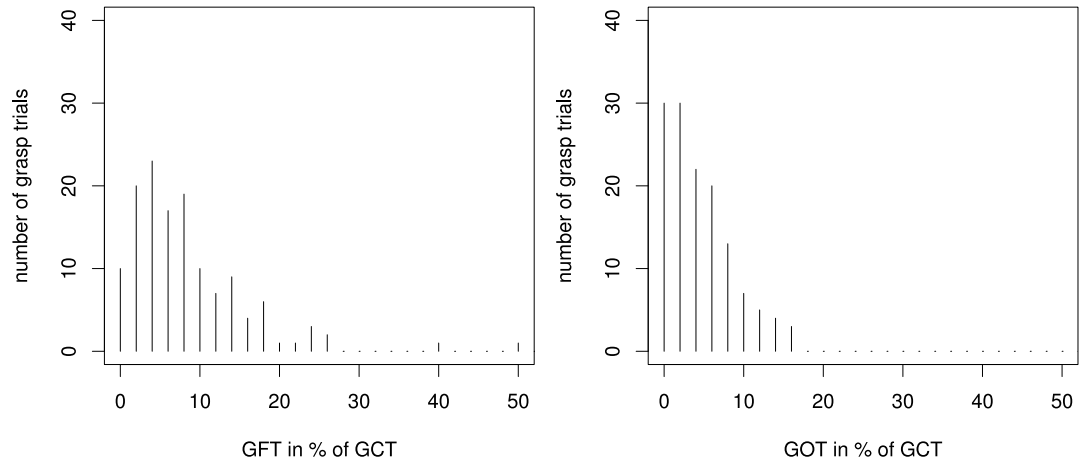

(c) GFT%; tea light.

(d) GOT%; tea light.

Figure 5.6: Distribution of grasp forming time (GFT) and grasp opposing time (GOT) in percentage of grasp closure time (GCT) over all trials of grasping the board marker and the tea light in a range of 2 percent. Means and standard deviations can be found in row "*all s.*" of Table 5.3.

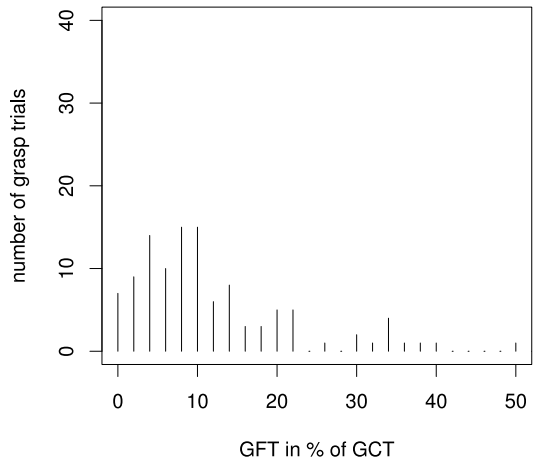

(a) GFT%; toy propeller.

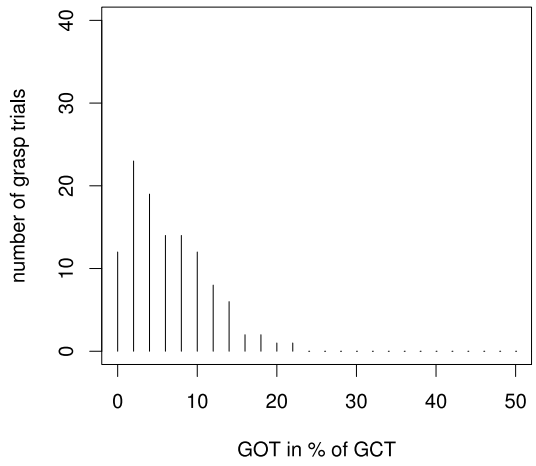

(b) GOT%; toy propeller.

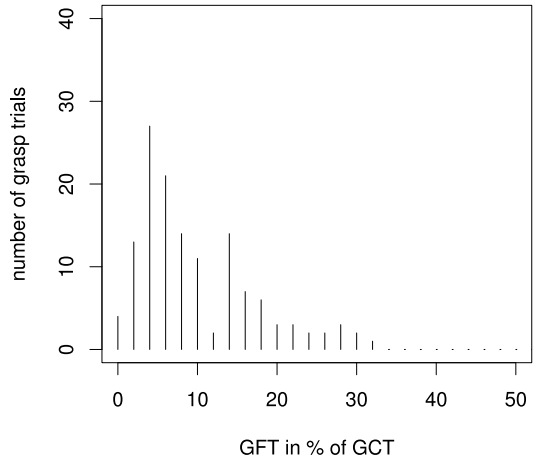

(c) GFT%; can.

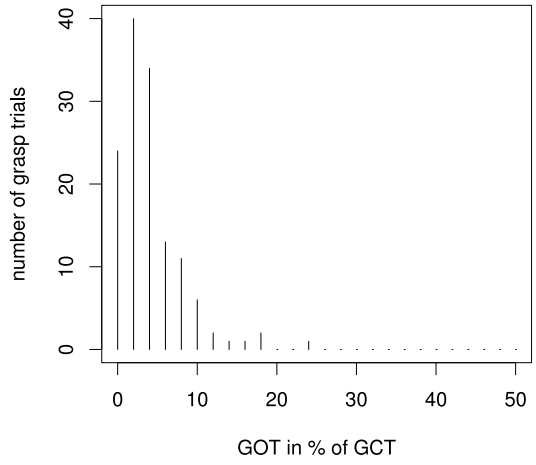

(d) GOT%; can.

Figure 5.7: Distribution of grasp forming time (GFT) and grasp opposing time (GOT) in percentage of grasp closure time (GCT) over all trials of grasping the toy propeller and the can in a range of 2 percent. Means and standard deviations can be found in row "*all s.*" of Table 5.3.

opposing (GOT). We conclude that both GFT% and GOT% prove that the subjects strove for contact simultaneity when closing the grasp.

ANOVA over all objects showed significant differences in GFT% ($F(3, 57) = 5,572$, $p = 0.002$) as well as in GOT% ($F(3, 57) = 2,849$, $p = 0.045$). Two-sample t-tests of the toy propeller versus each of the other three objects revealed that both GFT% ($p = 0.004$ against board marker, $p < 0.001$ against tea light, $p = 0.01$ against can) and GOT% ($p = 0.025$ against board marker, $p = 0.004$ against tea light, $p < 0.001$ against can) are larger. These results gives evidence that an increase of grasping difficulty (here caused by an enhanced complexity of the object's shape) leads to a decrease of contact simultaneity. This potential interrelationship was outside the scope of this experiment, but it is worth further investigation.

When comparing the distributions of GFT% between the different objects (see Figures 5.6 and 5.7), it can be seen that there are greater deviations from the mean in the case of the toy propeller. This goes along with the observation that some subjects in some grasp trials supported the grasp by one or two additional fingers, after they applied a grasp with the thumb and one or two opposing fin-

gers. Therefore, GOT% in some cases is much smaller than GFT% (compare GFT% and GOT% in column “toy propeller” of Table 5.3). This indicates that GOT% is a more reliable measure for contact simultaneity than GFT%.

5.6 Conclusions

In this experiment we proved that (i) humans strive for contact simultaneity when grasping objects. The values for the grasp forming time (GFT) and the grasp opposing time (GOT) are very small, and we found that there is a training effect towards even smaller values. That means that humans optimise their grasp process and achieve more contact simultaneity when repeatedly grasping the same objects. The analysis of GFT% and GOT% supports the conclusion that humans strive for contact simultaneity. We found that on average less than five percent of the grasp closure time (GCT) is required for establishing opposing contacts.

Contact simultaneity implies that a ”both” contact strategy is applied, which is performed in 93 percent of all grasp trials when using a rather short time frame of $t_f = 0.1s$. Besides striving for contact simultaneity, humans prefer a ”finger drag” to a ”thumb push” strategy. This result is independent of the value specifying t_f, and it is reasonable because if an object is moved before being grasped, the fingers can form a cage-like posture preventing the object from being moved away. This cannot be accomplished by the pushing thumb itself.

We found that (ii) the relative measures, GFT% and GOT%, are the most reliable measures for contact simultaneity. The assumption that the time required for grasp closure (GCT) affects GFT and GOT was proved. By relating GFT and GOT to GCT, the relative measures are independent of the velocity in which the grasps are performed. When comparing both measures, GOT% is the more reliable measure for contact simultaneity.

These findings support the assumption that the standard grasps of our robot hands can be optimised by a strategy that leads to more contact simultaneity. This topic is addressed in the following chapter. For measuring the success of this strategy in terms of improving contact simultaneity, we use the proposed measures for contact simultaneity, GFT% and GOT%, as described in Section 7.2.1.

6 Grasp Optimisation in Simulation

The five standard grasps realising the grasp types we propose for grasping a variety of everyday objects lead to greater or lesser success when grasping different objects as examined in Section 4.2.2. For optimising a grasp for grasping a particular object, we developed a strategy consisting of two steps. The first step optimises the pre-grasp posture for contact simultaneity, and the second step optimises the target grasp posture for the optimal trajectory of the thumb closure movement.

Most robot hand setups existing do not possess suitable sensors for performing these optimisation steps. Especially the second step can hardly be realised in real robot systems because it requires a value to determine the quality of a grasp being executed. Additionally, this step needs around 1000 grasp trials for optimising one target grasp posture. Thus, for optimisation, we use a physics-based grasp simulator containing an exact mapping of the robot hands, the benchmark objects, and the grasp types realised.

Section 6.1 describes this ”grasping world” and discusses the limitations of the simulator which have to be considered for determining a measure for grasp evaluation. To this end, we propose a grasp stability measure based on the grasp quality measure as discussed in Section 6.2. The grasp stability measure is used in the optimisation strategy we propose in Section 6.3.

6.1 The Simulated Grasping World

For grasp simulation, the physics-based toolkit Vortex [CMLabs, 2006] was employed within the graphical simulation environment Neo/NST [Ritter]. Vortex is a development platform that provides libraries for rigid-body dynamics, collision detection and response, and contact generation. It is based on Newtonian physics and realises accurate object motion and interaction capabilities. By adding features that incorporate static and dynamic friction for contacts and a quality measure based on the grasp wrench polytope [Haschke et al., 2005], we have a tool at our disposal that is suitable for analysing simulated grasping.

Within this framework 42 scenes (21 benchmark objects for grasping with both the TUM Hand and the Shadow Hand) have been created that can individually be loaded and processed by a single circuit of Neo-units. Figure 6.1 shows the control window in which one of these scenes is visualised. Directions and amplitudes of the acting forces are illustrated by blue lines and the finger contacts with the object are represented by friction cones. Green cones display static friction and are wider than the blue dynamic friction cones because the frictional coefficients are always larger in the case of static friction. The three-dimensional force part of the corresponding grasp wrench space (Figure 6.1 bottom right) is wider to its right side because of the occurrence of a static friction cone at the left side of the voltage tester (for more explanation see Section 6.2). Some of the most useful features that can be activated by buttons and sliders of the control window are:

- execution of all pre-grasps and all grasps,
- execution of both optimisation steps,
- engaging the stability calculation,

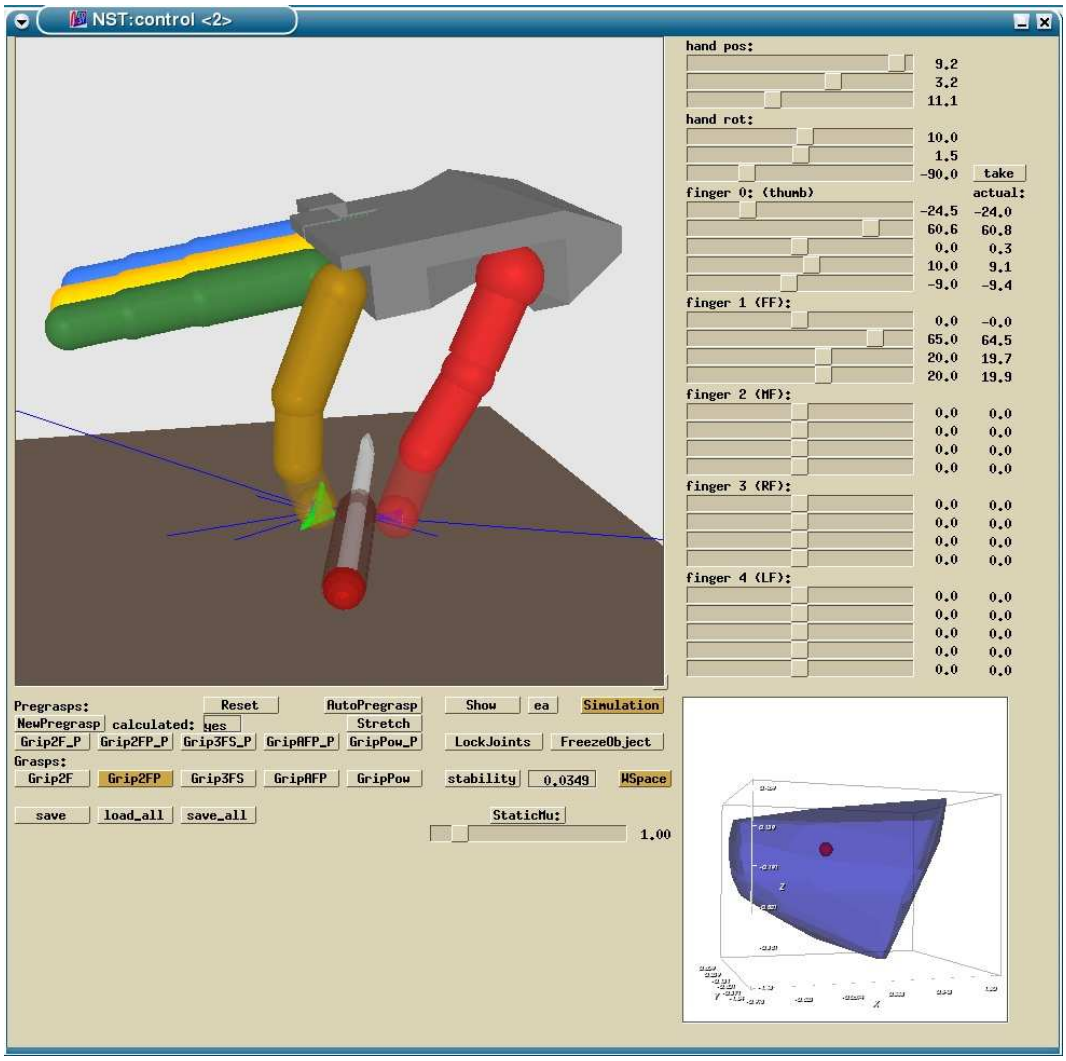

Figure 6.1: Control window showing the Shadow Hand grasping the voltage tester (one out of 42 scenes).

- control of position and orientation of the hand,
- control of all joint angles (number depends on hand type),
- alteration of friction coefficients.

In each scene of this Neo-circuit, the pre-grasps and the target grasps of all standard grasps implemented can be optimised for grasping the object loaded. Each grasp trial starts with the robot hand being in grasp position and in pre-grasp posture, before the target grasp is applied. It is not required to model the robot arm of each setup because only the manipulation component of the model of robot grasping proposed in Section 2.2.4 has to be considered for grasp optimisation.

Before the mapping of robot hands, objects, and grasps is described, in the following the contact modelling and limitations of the simulator are discussed.

Contact Modelling

Contacts between idealised rigid objects can be divided into point contacts, edge contacts, and face contacts [Nguyen, 1988]. For simplification, edge contacts are usually modelled as two point

contacts located at the ends of the edge, and face contacts are handled as point contacts at the vertices of the face [Ferrari and Canny, 1992].

Contacts between fingers and objects in real world, however, are not discrete points but extend over a certain area. How to model and handle the realistic distribution of forces and torques applied over this area is an open question. Especially in human grasping these contact areas are large because the soft finger surface deforms and assumes the local shape of the grasped object at the contact point. First approaches to simulate realistic grasping considering soft fingers exist [Ciocarlie et al., 2005], but like other current simulators that provide contact description (e.g. Miller and Allen [2004]), contacts in Vortex are simulated on the basis of point contacts with friction.

This contact modelling is, on the one hand, necessarily coarse. On the other hand, it is accurate enough for robot grasping, as the contact surfaces of robot fingers are at most as soft as rubber and thus establish much smaller contact areas than human fingers.

A real contact between a finger and an object is not necessarily modelled as a single point contact. The number of point contacts occurring in one simulated contact area depends on the surface models and the effects of interpenetration as described in the following sections.

6.1.1 Limitations of the Simulator

The simulator is designed for real-time, interactive simulation, and visualisation. Thus, it has to balance between accuracy and speed. This balance means that limitations concerning physics calculation, details of scene modelling, and collision detection and handling are inevitable. These limitations have to be considered when analysing grasping in simulation.

A central problem in modelling and simulating real scenes is the discretisation of space and time. This discretisation leads to the undesired effect of object interpenetration. When grasping an object in simulation, the fingers move by discrete distances in separate time steps rather than moving continuously. Therefore, in the time step before a finger touches an object no collision can be detected, while in the next time step the finger and the object already interpenetrate. This effect has to be counteracted, and for this purpose, a restoring force is simulated until the penetration is reversed. As an undesired result, it may happen that the contact is lost in the next time step. This leads to an unrealistic object motion between the fingers, additionally supported by the fact that the amount of the restoring force is dependent on the depth of interpenetration. Resultant forces acting on fingers and object may vary greatly between two time steps.

Another challenge arising from the interpenetration effect is that an unpredictable number of point contacts may occur between finger and object. Even when using simple geometrical models for which only one point contact would be expected, the interpenetration enlarges the contact area and that may result in several point contacts at different locations.

The third problem to be faced is that restrictions in the computational accuracy lead to rounding errors in complex calculations of collision detection. As a result, it happens that contacts switch between static and dynamic friction. This effect is also depicted in Figure 6.1 which shows a static friction cone to the left hand side of the object, whereas on the right hand side only a smaller dynamic friction cone can be seen. Although in the time step illustrated the motions of fingers and object have ended already, in the next simulation step dynamic contacts could switch to static ones and vice versa [Steuwer, 2003].

To summarise, problems to be dealt with are unrealistic object motion between fingers, unpredictable numbers of contact points at different locations, and spontaneous switches between static

and dynamic contacts. Therefore, a simulation with absolute realistic behaviour is not attainable. To achieve a sufficiently realistic simulation of robot hand grasping being exact enough for realising optimised grasps that lead to more grasp success when applied with the real hands, a careful modelling of objects, hands, and grasps and a precise adjustment of the parameters required are crucial as discussed in the following section.

6.1.2 Mapping of Objects and Hands

For mapping a benchmark object or a robot hand into simulation, a model consisting of a single or several objects has to be created. In Vortex, basic objects like spheres, boxes, or cylinders can be defined. More complex shaped objects comprise a surface model consisting of a mesh of polygons. When creating an object (for instance, with a 3D object design tool), the level of detail of the surface model has to be carefully balanced because the more polygons a surface model consists of, the higher are the computing demands on the collision detection and handling.

As discussed in Section 6.1.1, the interpenetration of two surfaces enlarges the contact area. The polygons of a surface mesh are handled as single surfaces and that may result in several point contacts at the boundary of the contact area.

Modelling objects

For each of the 21 benchmark objects described in Section 4.1, a model was created. In designing these models, we considered the characteristics of the benchmark objects such as dimensions, mass, centre of mass, and coefficients of static and dynamic friction between the real object and the desktop surface (as given in Table 4.1). By balancing between detail and a low number of surface polygons, we created realistic models (pictured in Figure 4.1) that do not overstrain the collision handler.

Form-variable benchmark objects are modelled as simple boxes. The reason is that an exact mapping of form-variable objects consisting of rigid components, like the folding rule or the bunch of keys, is hardly feasible and would result in high computing costs. To model the form-varying properties of a paper ball is not realisable with rigid objects at all. But boxes of the dimensions determined as described in Section 4.1.2 are sufficiently exact models of these benchmark objects for determining optimised grasps as evaluated in Chapter 7.

Modelling hands

When modelling the robot hands, an exact mapping of the parts of the palm and fingers in their dimensions, in their locations, and in their orientations is crucial for a realistic grasp simulation. Even small discrepancies between the modelled hand and the real hand can lead to a totally different grasp behaviour. The surface models of the hand parts were designed by balancing the detailedness of the models.

Modelling the hydraulic piston drive of the TUM Hand, or the muscle dynamics and the complex routing of tendons inside the Shadow Hand, is infeasible. However, for realising the closure behaviour of the fingers, it is sufficient to use simple hinge models for the finger joints and determine appropriate parameter values for the maximum joint velocities and the maximum joint forces.

To achieve realistic interaction between the fingers of a robot hand and the target object, suitable coefficients of friction have to be determined between the material of the hand's contact areas and

the material of the object. The surfaces of the fingertips, in the case of the TUM Hand, and the volar surfaces of the fingers, in the case of the Shadow Hand, show similar properties. Frictional coefficients between these surface materials and the rubbery material of the eraser (object no. 19) were determined to be $\mu = 2.0$ (static friction) and $\nu = 0.6$ (dynamic friction). Frictional coefficients between the surface materials of the hands' contact areas and the materials of all other objects were determined to be $\mu = 1.0$ and $\nu = 0.3$. For optimising grasps of objects that can hardly be grasped, these frictional coefficients can be enlarged as described in Chapter 7.

6.1.3 Mapping of Grasps

With an exact model of a real robot hand and its finger closure behaviour in simulation, realistic mappings of the standard grasps developed in Section 3.2.3 are implemented. To this end, the same parameters identified for the real standard grasps are used in simulation. The approach distance d (see Section 3.2.1) is not required because the transportation component of the model of robot grasping can be neglected for grasp optimisation. The position $\boldsymbol{p}$ and the orientation $\boldsymbol{o}$, identified for the standard grasps, are presented in Table 3.1, and the pre-grasps and the target grasps are listed in Tables A.1 and A.2. The mappings of the grasps, with both simulated hands in pre-grasp posture and in target grasp posture, are shown in Figure 6.2.

6.2 Measures for Static and Dynamic Grasps

For grasp evaluation, comparison of different grasps, and to provide the grasp optimisation algorithm with a scalar fitness value, a measure is required that provides a reliable rating of the grasp. Based on the definitions of force closure and the grasp wrench space given in Section 6.2.1 (for detailed definitions see Nguyen [1988] or Ferrari and Canny [1992]), we determine a quality measure that implies forces in task direction as explained in Section 6.2.2.

While the grasp quality measure is calculated for a static set of contact points, the grasp stability measure we propose considers the grasping process as a dynamic system. Determined in simulation, the grasp stability is a reliable measure for estimating the grasp success when applying the same grasp with the real hand. Section 6.2.3 compares the traditional definitions of grasp stability with ours, before Section 6.2.4 explains our definition in detail.

6.2.1 Force Closure and the Grasp Wrench Space

A set of contact points on the surface of an object is termed *grasp configuration*. A grasp configuration is called *force closure* if forces $\boldsymbol{f}_{c_i}$ at the contact points c_i (modelled as point contacts with friction, see Section 6.1) can be applied on the object such that any external force acting on the object can be balanced. Some grasp configurations can be better than others in the sense that they can balance larger external forces by applying the same amount of finger forces.

A force applied at a contact point results in a force $\boldsymbol{f} \in \mathbb{R}^3$ and a torque $\boldsymbol{\tau} \in \mathbb{R}^3$ acting on the objects centre of mass. They can be represented as a *wrench* $\boldsymbol{w}$ in the 6-dimensional space:

$$\boldsymbol{w} = \begin{pmatrix} \boldsymbol{f} \\ \boldsymbol{\tau} \end{pmatrix} \in \mathbb{R}^6 \tag{6.1}$$

To determine the amount of every external force that can be resisted by a grasp configuration, the *grasp wrench space* $\mathcal{W}$ has to be examined. $\mathcal{W}$ determines the set of wrenches that could be

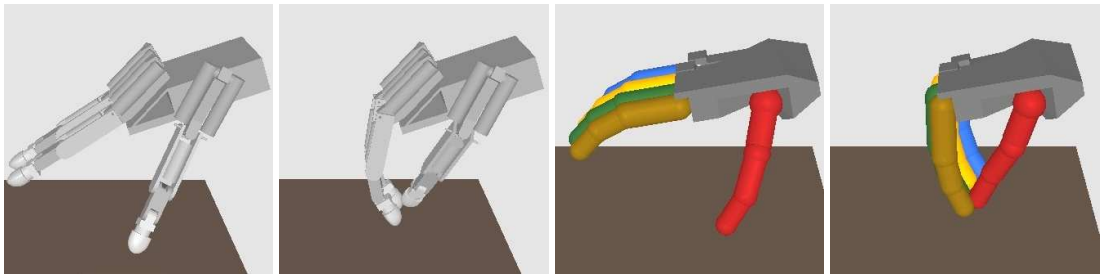

(a) All finger precision grasp g_1.

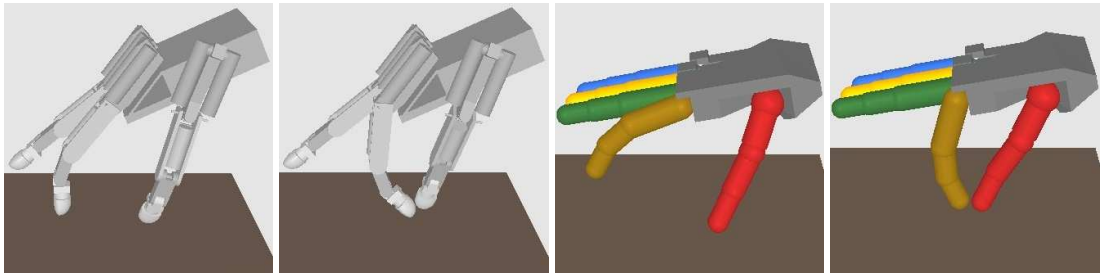

(b) Two finger precision grasp g_2.

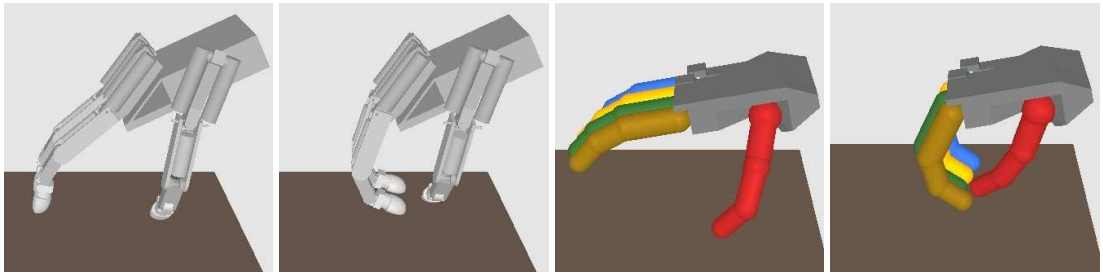

(c) Power grasp g_3.

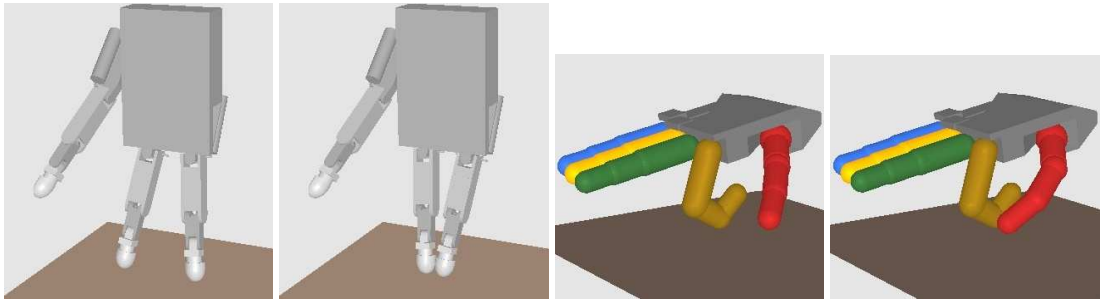

(d) Two finger pinch grasp g_4.

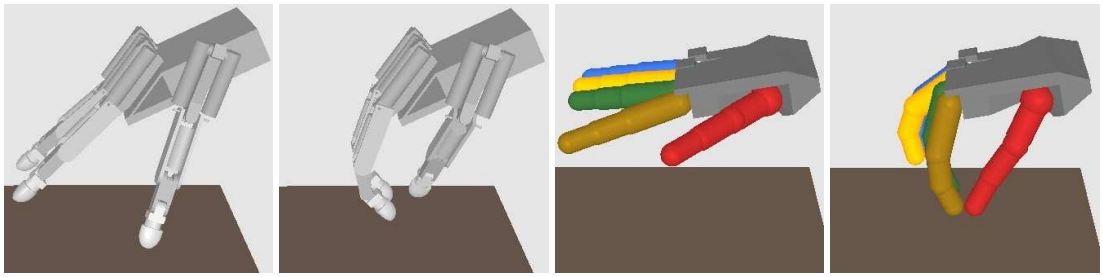

(e) Three finger special grasp g_5.

Figure 6.2: Pre-grasp and target grasp postures of the TUM Hand and the Shadow Hand.

exerted on the target object and can be approximated by limiting the contact normal forces $\boldsymbol{f}_{c_i}^{\perp}$ as described below. The set of forces $\boldsymbol{f}_{c_i}$ applicable at a contact c_i is defined by its friction cone depicted in Figure 6.3.

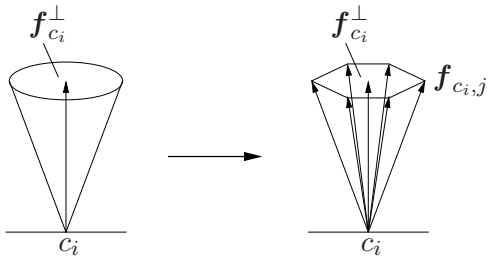

Figure 6.3: Complexity reduction of a contact's friction cone.

Each contact force $\boldsymbol{f}_{c_i}$ can be decomposed in a linear combination of force vectors lying on the boundary of the contact's friction cone. To reduce complexity in computing the wrench space, the friction cones are approximated by a finite number m of vectors $\boldsymbol{f}_{c_i,j}$ (see Figure 6.3).

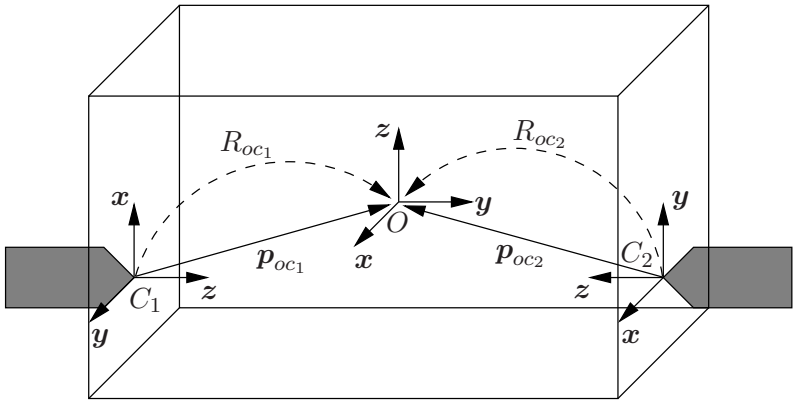

Figure 6.4: Transformation of contact frames C_i into the object frame O.

By transforming the boundary force vectors $\boldsymbol{f}_{c_i,j}$ of each contact frame C_i into the coordinate frame of the object O (origin o equals the centre of mass), the corresponding wrenches $\boldsymbol{w}_{c_i,j}^{o}$ acting on the object can be calculated (compare Figure 6.4):

$$\boldsymbol{w}_{c_i,j}^{o} = \begin{pmatrix} R_{oc_i}^{T} \boldsymbol{f}_{c_i,j} \\ R_{oc_i}^{T}(-\boldsymbol{p}_{oc_i} \times \boldsymbol{f}_{c_i,j}) \end{pmatrix} \tag{6.2}$$

R_{oc_i} is the rotation matrix and $\boldsymbol{p}_{oc_i}$ the translation vector to transform the contact frame C_i into the object frame O. A convex hull of the wrenches $\boldsymbol{w}_{c_i,j}^{o}$ approximates the grasp wrench space $\mathcal{W}$, whereas the contact forces are typically limited to unity with respect to their normal component $\boldsymbol{f}_{c_i}^{\perp}$, using the metrics L_∞ or L_1 [Ferrari and Canny, 1992]:

$$\begin{array}{lll} 1. \quad \|\boldsymbol{f}_c\|_\infty = \max_i |\boldsymbol{f}_{c_i}^{\perp}| \leq 1 & \mathcal{W}_{L_\infty} = co\left(\bigoplus_{i=1}^{n} \left\{\boldsymbol{w}_{c_i,1}^{o}, \ldots, \boldsymbol{w}_{c_i,m}^{o}\right\}\right) & \\ 2. \quad \|\boldsymbol{f}_c\|_1 = \sum_i |\boldsymbol{f}_{c_i}^{\perp}| \leq 1 & \mathcal{W}_{L_1} = co\left(\bigcup_{i=1}^{n} \left\{\boldsymbol{w}_{c_i,1}^{o}, \ldots, \boldsymbol{w}_{c_i,m}^{o}\right\}\right) & \mathcal{W}_{L_1} \subseteq \mathcal{W}_{L_\infty} \end{array} \tag{6.3}$$

When utilising the L_∞ metric (Equation 6.3.1) the applied contact forces are individually limited. That leads to an enlargement of the convex hull $\mathcal{W}_{L_\infty}$ with each contact. Because of the limitations in grasp simulation (see Section 6.1.1), the number of contacts at each contact region depends on the surface materials used. Additionally, it varies between different simulation time steps. The differences in the enlargement of the approximated grasp wrench space between two objects or consecutive time steps can be large when using the L_∞ metric. This cannot be a reliable basis for a grasp quality measure. Thus we use the L_1 metric (Equation 6.3.2), which limits the sum of all contact normal forces $\boldsymbol{f}_{c_i}^{\perp}$ to one and thus limits the potential enlargement of the convex hull $\mathcal{W}_{L_1}$.

$\mathcal{W}_{L_1}$ has a geometrical meaning to the force closure requirement of the grasp configuration. Because it represents all wrenches that could be exerted on the target object by the grasp configuration, the requirement of being force closure is fulfilled, if the origin is located inside the convex hull and not on its surface (compare Figures 6.5b and 6.5c). In this case, wrenches on the object in any direction can be exerted, and thus disturbances from any direction can be resisted.

6.2.2 Grasp Quality Measure

We determine the *grasp quality measure* by taking the amplitude of the most unfortunate disturbance wrench that could just be resisted by the grasp configuration, given limited contact forces as defined in Section 6.2.1. This amplitude corresponds to the radius α of the largest hyper sphere centred in the origin o that just could be completely enclosed by the convex hull. In other words, it corresponds to minimal distance of the origin o to the surface of the convex hull.

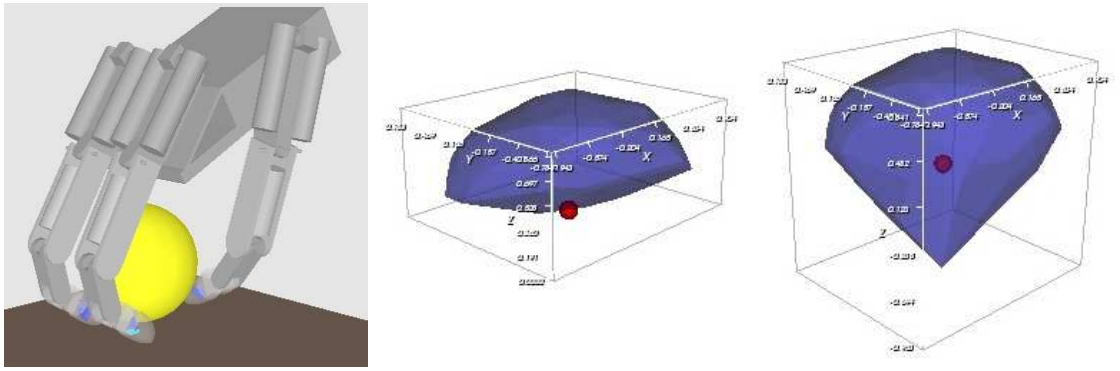

(a) Fingertips touch the ball at its bottom area.

(b) Grasp is not force closure ($\alpha = 0$) because the object origin o is on the surface of the convex hull.

(c) Grasp is force closure ($\alpha = 0.178$) when considering a force in direction of gravity.

Figure 6.5: Expanding $\mathcal{W}_{L_1}$ with a force opposed to "task direction" leads to an enlarged convex hull and an appropriate evaluation of grasp quality. The three-dimensional force component of the convex hull, which corresponds to the grasp configuration shown in (a), is illustrated in (b) and (c).

The calculation of the approximated L_1 grasp wrench space $\mathcal{W}_{L_1}$ considering equations (6.2) and (6.3.2) results in a discrete set of six-dimensional points. $\mathcal{W}_{L_1}$ can easily be expanded by further points in order to calculate a quality measure that considers desired directions of forces to be applied on the object [Haschke et al., 2005]. For example, the task of bouncing a ball requires an application of force in the direction towards the ground. At the time when the hand touches the ball and the fingers exert that force, this "interaction grasp" is not force closure although it could be quite appropriate for bouncing the ball. The task of grasping an object from a desktop implies a force in opposite direction (against gravity). Grasps that are not force closure though accom-

plishing this task (for example, when the fingertips touch a ball at its bottom area as shown in Figure 6.5a) are considered to be appropriate for this task if $\mathcal{W}_{L_1}$ is expanded by a force in direction of gravity (compare Figures 6.5b and 6.5c). This is done by adding the point $(0, 0, -1, 0, 0, 0)^T$ which corresponds to the wrench of a frictionless point contact applying the desired force whereas causing no torque on the object's centre of mass [Steffen, 2005]. The resulting expanded set is

$$\mathcal{W}_{L_1,grav} = co\left(\bigcup_{i=1}^{n}\left\{\boldsymbol{w}^o_{c_i,1},\ldots,\boldsymbol{w}^o_{c_i,m}\right\} \cup \left\{(0,0,-1,0,0,0)^T\right\}\right) . \tag{6.4}$$

Grasp Quality of Different Simulation Steps

The simulation of a grasp comprises the grasp closure phase and the stabilisation phase (compare Section 2.2.4). Starting from the hand in pre-grasp posture in the first simulation time step s_1, the target grasp is applied during the following simulation steps ($s_1, s_2, \ldots, s_g, s_{g+1}, \ldots$). Time step s_g is the first step, in which all grasping fingers touch the target object.

A grasp quality value α based upon the expanded grasp wrench space $\mathcal{W}_{L_1,grav}$ can be calculated in every simulation step s_{g+k} ($k = 0, 1, \ldots$) in which the grasping fingers touch the object. The suitability of different steps s_{g+k} for providing a reliable rating of the grasp is discussed in the following:

- First step in which all grasping fingers have contact s_g:
  In the first step s_g when all grasping fingers have contact with the object, the magnitude of interpenetration (as described in Section 6.1.1) can be low or high. The resulting number and locations of contact points are casual. Thus, a grasp quality value α calculated in s_g is most unreliable.

- Step in which finger motions stop:
  Two different reasons lead to the stop of the finger motions. One is that the target object prevent the fingers from further motion, and another is that the fingers reach their target grasp posture. The latter alternative points to the fact that the object is lost. In some cases, this corresponds to realistic grasp behaviour, but many objects are lost in simulation after being grasped properly. Again, the reason is the effect of interpenetration. It leads to too large counterforce acting on the object, by which the object bounces between the fingers. Because this is an unrealistic behaviour, the step in which finger motions stop does not provide a reliable grasp rating.

- Second s_{g+1} or later step in which all grasping fingers have contact:
  In deciding which simulation time step s_{g+k} best represents the quality of the grasp, one problem is that the quality value α in subsequent steps (s_{g+k}, s_{g+k+1}) may vary largely. This effect can be observed particularly when an object is touched at one or more edges. In these cases, the orientation of contact friction cones can change decisively (compare Figures 6.6a and 6.6b). If s_{g+1} or a later simulation step in which all grasping fingers have contact is used for calculating the grasp quality, this value could be too low or too high, and thus it would not be reliable.

The grasp quality measure α evaluates a static grasp configuration. But the discussion of the quality values calculated in different simulation steps s_{g+k} shows that no single grasp configuration is appropriate for providing a reliable rating of the grasp. Because grasping with a robot hand is a dynamic process, an object that is grasped in one step can be lost in the next. Although a grasp

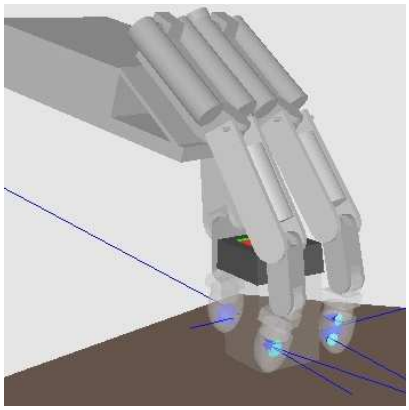

(a) Preceding simulation step with grasp quality value $\alpha = 0.126543$.

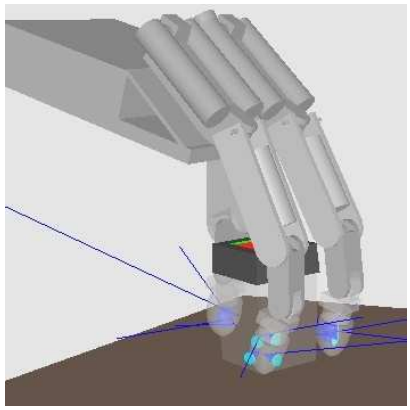

(b) Subsequent simulation step with grasp quality value $\alpha = 0.258619$.

Figure 6.6: Index and middle finger touch edges of the sharpener. In subsequent simulation steps (s_{g+k}, s_{g+k+1}) slightly different contact locations lead to large changes in friction cone orientations and in quality values α.

configuration with a quality value $\alpha > 0$ can theoretically resist disturbances from any direction, the wrenches actually applied can move the object into a grasp configuration with a quality value $\alpha = 0$. Therefore, a grasp evaluation measure is needed that considers the grasp as a dynamic system and evaluates its "stability".

6.2.3 Grasp Stability

The concept of stable grasping was introduced by Hanafusa and Asada [1977] and extended by Nguyen [1986]. The criteria for *stability* are that (1) the object be in static equilibrium; and (2) the grasp configuration be force closure. An object is in *equilibrium* if the sum of the wrenches acting on the object is zero. The fulfilment of the theoretical force closure criterion is not sufficient when a dynamic system like a robot hand has to realise the grasp configuration. The forces being actually applied and forces the manipulator is able to apply have to be considered. Thus, Fearing [1986] defines two more requirements for grasp stability: (2a) the direction of the applied forces must be within the friction cones; and (2b) it should be possible to increase the magnitude of the grasping applied forces to prevent any displacement due to an arbitrary disturbance force.

In our grasping simulation, the grasps can be tested on the force requirements (2), (2a), and (2b). To fulfil criteria (1), the target object must not move between the grasping fingers. Because of inevitable interpenetration (see Section 6.1.1), this can hardly be fulfilled when grasping in simulation. To achieve this goal, we tried to fixate the object and the fingers during the grasp execution in three different ways:

- Fixating fingers when touching the object:
  To get rid of the object motion caused by interpenetration effects, one idea was to stop the fingers from further motion when all grasping fingers touch the object (in step s_g). This inhibits the interaction of motion between the fingers and the object, and it was expected that the object stops its motion too. But test trials showed that although counteracting forces were lowered by stopping the fingers, residual penetration with the fingers caused the object to "bounce" between them. Additionally, the effect that contacts switch between static and dynamic friction arised (also described in Section 6.1.1), leading to very different grasp quality values between subsequent simulation steps.

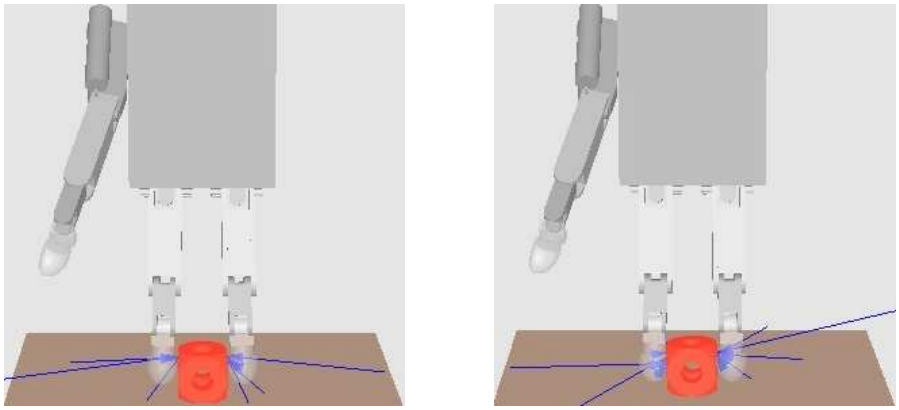

(a) Fingers touch the fixated object simultaneously.

(b) More realistic grasp behaviour is achieved with an un-fixated object.

Figure 6.7: Comparison of fixated and freely moving cube while grasping it with the TUM Hand in simulation.

- Fixating object while grasping:
  The first criteria for stability is definitely fulfilled if the object is fixated during grasp simulation. But this method does not lead to a realistic grasp process for two reasons. In most cases when grasping with both real robot hands, the object is pushed by the thumb (*thumb push strategy*) or dragged by one or more other fingers (*finger drag strategy*) before all grasping fingers have contact. These effects are inhibited when fixating the object in the simulated world. The second reason is that even in the case of simultaneous finger contact, the object is moved between the fingers before a ”stable” grasp is established. This behaviour can be observed when grasping with a real robot hand and also when grasping an un-fixated object in simulation. Thus, fixating the object leads to an inaccurate grasp evaluation (compare Figures 6.7a and 6.7b).

- Fixating object and fingers after contact:
  When stopping both the motion of the fingers and the motion of the object as soon as the grasping fingers have contact (in step s_g), the criteria (1) for stability is also fulfilled. But this stops the dynamic behaviour and the current distribution of point contacts is conserved. Because of the described interpenetration effects, especially in this simulation step s_g the distribution of contacts and forces is unreliable. Additionally, this static grasp does not shed light on the possibility of losing the object in a later simulation step.

None of the three approaches leads to a reliable grasp evaluation measure. But without fixating the object or the fingers in simulation, grasp stability cannot be achieved according to its traditional definitions. These definitions divide grasp configurations into two categories: stable or not stable. Each definition requires that forces applied at the contact points can be infinitely large to resist arbitrary disturbances. When considering the limited forces a real manipulator can apply, no grasp configuration can be stable. Therefore, we propose a new definition of grasp stability that considers the manipulator used. Instead of distinguishing two categories, it provides a measure for the probability that a target object is grasped successfully when applying the grasp.

6.2.4 New Approach to a Grasp Stability Measure

We define ”grasp stability” as a measure based on the grasp quality values calculated in appropriate simulation steps. Because single simulation steps cannot provide a basis for a realistic grasp

evaluation value (as discussed in Section 6.2.2), the *grasp stability measure* σ is defined as an average over the grasp quality values α calculated in different simulation steps s.

Before giving the formal definition of the grasp stability measure, this section describes the major considerations of which simulation steps s lead to the most reliable grasp rating.

Averaging over a Number N of Subsequent Grasp Quality Values

In each simulation step s_{g+k}, a quality value α of the actual grasp configuration can be calculated. Subsequent grasp quality values are quality values calculated in subsequent simulation steps in which all grasping fingers touch the object.

One approach to acquiring quality values α that are not necessarily calculated in subsequent simulation steps is choosing only quality values that are larger than zero ($\alpha > 0$). But this approach does not lead to a reliable grasp stability measure σ in any case. For example, if a grasp is at the limit to force closure (for instance, the TUM Hand grasping the board marker after the first optimisation step, see Section 7.1.1), subsequent single quality values α may alter between zero and a value larger than zero. If only simulation steps s_{g+k} are considered in which $\alpha > 0$, the average over a number of these quality values may be as large as that of a grasp in which grasp configurations of subsequent simulation steps are force closure. But the chance of applying a successful grasp with a real hand is higher in the latter case.

Thus, we propose to average over a number N of subsequent grasp quality values, independent whether $\alpha > 0$ or not.

The First Grasp Quality Value α_1 to be Considered

Besides determining the number N of subsequent grasp quality values α for averaging, the first simulation step to be considered has to be determined. Considering the first step in which $\alpha > 0$, leads to an unreliable grasp stability σ because of two reasons.

On the one hand, especially in unreliable two finger grasps, it may happen that the first quality values are zero, although the grasping fingers touch the object already. This fact is neglected, if the first quality value α_1 to be considered is determined by a value larger than zero, and the resulting stability value σ would be too large. On the other hand, the grasp stability value σ could be too small in case a grasp type t with more than two fingers is applied. A quality value larger than zero ($\alpha > 0$) can be achieved with the first two fingers touching the object (because of interpenetration more than two point contacts can occur). This value can be very small compared to the grasp quality value calculated when all grasping fingers touch the object.[1]

In summary, for determining the first simulation step s to be considered for calculating the grasp stability value σ, the height of the grasp quality α is not taken into account, and α calculated in simulation step s_g is most unreliable as discussed in Section 6.2.2. Therefore, we propose to define the second step s_{g+1} in which all grasping fingers touch the object as the first step considered for calculating the grasp stability value σ. The simulation steps s_g as well as s_{g+1} can be identified by counting the fingers having contact with the object and comparing to the number f of fingers that should touch the object. This number f has to be determined for each grasp type t and can be object-specific in the case of grasping with the Shadow Hand.

[1] For instance, when grasping the remote control with the Shadow Hand by utilising the standard grasp, the resulting stability value ($\sigma = 0.076190$) is approximately one-quarter of that value ($\sigma = 0.284305$) achieved when applying the more reliable approach we use.

Number of Grasping Fingers f

Before a grasp is applied with a five-fingered hand, like the Shadow Hand, it has to be determined how many fingers f will actually touch the object. If a two finger grasp is executed, both grasping fingers have to touch the object because at least two fingers are needed to provide opposing forces ($f = 2$). But in the cases of the power or the all finger precision grasp, some target objects are only touched by three of four instead of five fingers. For example, a small-sized object, like the golf ball, is only touched by the thumb, the middle finger, and the ring finger ($f = 3$).

When grasping with the three-fingered TUM Hand, no object-specific distinction has to be made. In the case of applying a two finger grasp, again both fingers have to contact the object for a stable grasp ($f = 2$). When utilising other grasp types, each benchmark object is touched by all three fingers of the TUM Hand ($f = 3$).

Phalanges to be Considered for Contacts

For determining the quality values α to be considered, the number of fingers having a ”valid” contact with the object have to be counted. A contact is valid only if specific phalanges touch the object. The appropriate phalanges have to be identified for each grasp type t.

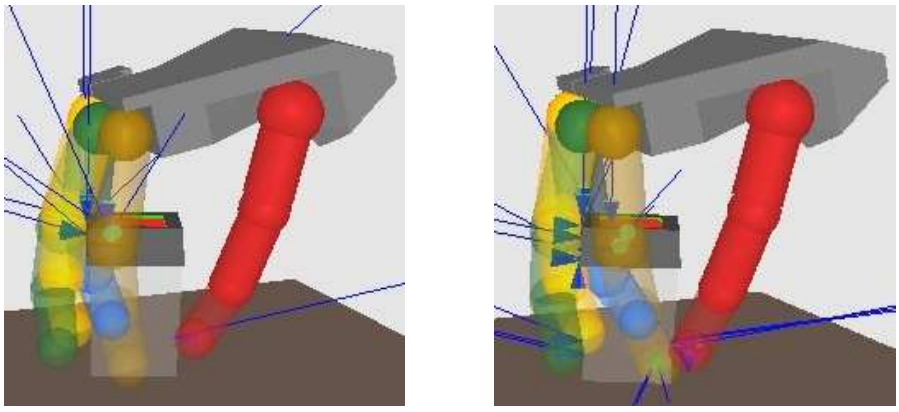

(a) Only the thumb touches the object with its tip first.

(b) The grasp closure phase ends when all fingertips touch the object.

Figure 6.8: Grasping the sharpener with the Shadow Hand is an example of a precision grasp in which the fingertips do not touch the object first.

The main characteristic of **precision grasps** (t_1, t_2, and t_5) is that the fingertips make contact with the object. Thus, a finger has a valid contact if the fingertip touches the object. Most benchmark objects are touched by the tips of the fingers first. But, for example, when grasping the sharpener with the Shadow Hand, the proximal phalanges touch the object before as shown in Figure 6.8a. This grasp configuration is not force closure, and a stability evaluation which does not consider which phalanges touch the object would result in a value $\sigma = 0$. But the grasp closure phase does not end before the fingertips touch the object (see Figure 6.8b), and this time step has to be the first to be considered for calculating the grasp stability σ.

When utilising a **power grasp** (t_3), usually more phalanges touch the object, and even the palm might be involved (for example, when grasping the tennis ball as depicted in Section 7.1.2). In contrast to precision grasps, contacts with the fingertips are not that important. Because the fingers

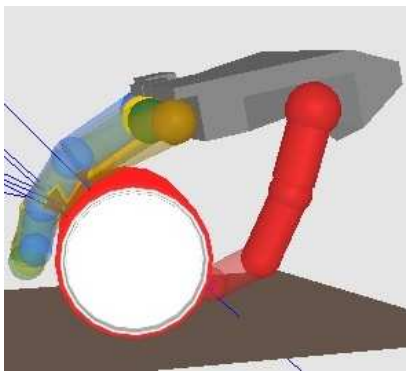

(a) Force closure grasp without fingertip contacts, except for the tip of the thumb.

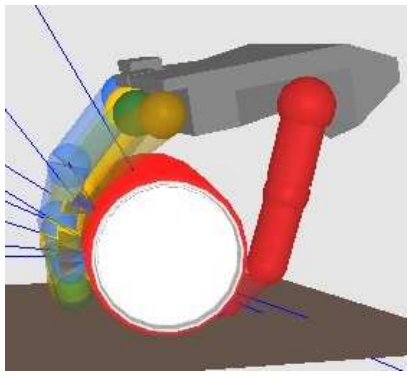

(b) Object is enclosed when all fingertips have contact.

Figure 6.9: Grasping the can with the Shadow Hand is an example of a power grasp in which the fingertips do not have first contact.

enclose the object, contacts with other phalanges can be sufficient for the grasp to be force closure (see Figure 6.9a). But the grasp closure phase does not end until all fingertips have contact, and the object is fully enclosed (see Figure 6.9b). Thus, contacts are valid if the object is touched by the fingertips.

The **pinch grasp** (t_4) of the TUM Hand essentially differs from the pinch grasp executed with the Shadow Hand. When determining the finger parts to be considered, we have to distinguish between the hands used. In the case of the TUM Hand, only the sides of the fingertips are suitable for touching the object because of the hardware restrictions described in Section 3.2.3. Thus, a contact is valid if the fingertip touches the object. But when applying the pinch grasp with the Shadow Hand, the target object is pressed by the thumb against the side of the flexed index finger. Because the sides of the middle and the proximal phalanges of the index finger may provide opposing forces, contacts with these phalanges, besides contacts of the tips of the thumb and the index finger, have to be considered when applying a pinch grasp with the Shadow Hand.

Determining the Number N of Quality Values α for Averaging

There are several requirements for the number N of quality values α for averaging when calculating the grasp stability σ: (i) Subsequent quality values may differ largely, for instance, when the object is grasped at an edge. To get a reliable grasp stability value σ, the number N of quality values must not be too low. (ii) A too large number N leads to a stability value σ being too low in cases when the object is lost after being grasped properly due to the effects of interpenetration (see Section 6.1.1). (iii) If the forces exerted by the fingers prevent the object from reaching equilibrium, the object is lost after some simulation steps. This realistic behaviour may or may not lead to success when grasping with the real hand. One example is the grasp of the light bulb with any of the robot hands used. This uncertainty has to be taken into account when calculating the grasp stability value σ. Thus, the number N of quality values α for averaging has to be large enough to consider some simulation steps after the object is lost and the grasp quality values are zero.

After gaining a lot of experience with applying different grasps in each of the 42 simulated scenes, we identified the number N of quality values α to be equal to 10. This number best considers all of the three requirements and results in a reliable grasp stability value σ.

Definition of the Grasp Stability Measure σ

Summing up all constraints for defining a reliable *grasp stability measure* σ described in this section, we propose the following definition.

The *grasp stability measure* σ of a grasp g is defined as

$$\sigma = \frac{\sum_{i=1}^{N} \alpha_i}{N} \qquad i = 1, \ldots, N\,, \tag{6.5}$$

where

$$N = 10\,.$$

The grasp quality values $\alpha_1, \ldots, \alpha_N$ are calculated in subsequent simulation time steps in which all grasping fingers have valid contacts. N is the number of subsequent grasp quality values for averaging.

The first simulation time step s_{g+1} to be considered for calculating the first grasp quality value α_1 is determined by the second step in which all grasping fingers have valid contacts.

A *valid contact*, is a contact with the distal phalanx, i.e. the fingertip, for all grasp types t and both robot hands with one exception: utilising a two finger pinch grasp t_4 with the Shadow Hand. In this case, a contact with every phalanx is a valid contact.

To identify whether all grasping fingers have valid contacts, the number of fingers having a valid contact is counted in each simulation step s. This number is compared to the number f of grasping fingers defined, which depends on the used grasp type with object-specific exceptions in the case of the Shadow Hand:

$f = 2$; for the types: two finger precision grasp t_2, and two finger pinch grasp t_4;
$f = 3$; for the type: three finger special grasp t_5;
$f = f_{max}$; for the types: all finger precision grasp t_1, and power grasp t_3;

where f_{max} is the maximal number of fingers. In the case of the TUM Hand, $f_{max} = 3$; in the case of the Shadow Hand, $f_{max} = 5$.

The object dependent exceptions when grasping with the Shadow Hand are:

$f = 4$; for the objects: adhesive tape, tennis ball, cup, tea light, bunch of keys;
$f = 3$; for the objects: sharpener, golf ball.

If the number M of quality values α that can be calculated before the target grasp posture is reached (the object is lost) is less than $N = 10$:

$$\sigma = \frac{\sum_{i=1}^{M} \alpha_i}{N} \qquad i = 1, \ldots, M\,. \tag{6.6}$$

6.3 The Optimisation Strategy

With our definition of grasp stability, we have a reliable measure to compare different grasps in simulation. Grasps with a larger stability value are supposed to lead to more successful trials when grasping with the real hand. In our approach to grasp optimisation, the stability value σ is used to evaluate the optimisation progress. We propose a strategy that optimises both the pre-grasp and the target grasp for a particular target object in two subsequent steps:

1. One/Two-shot learning of the pre-grasp.
2. Optimising the target grasp with an evolutionary algorithm.

The first optimisation step is based upon experiences made when evaluating the standard grasps developed for the TUM Hand as described in Section 4.2. We found that the longer it takes the last grasping finger to reach the target object, the smaller the number of successful grasp trials. This is primarily due to the displacement of the object by the fingers touching it first. Therefore, we suggest that a grasp can be optimised by reaching finger contact simultaneity. This assumption is supported by the results of our grasp experiment described in Section 5.6.

The second step optimises the target grasp. In the case of the Shadow Hand, the target grasp comprises 18 different joint angles. Optimising all joint angles together implies a search in a space of 18 dimensions. Whatever optimisation method is applied, a search in such high-dimensional space would be very time consuming, and whether the global optimum could be found is questionable. Many local minima can be expected which would not be appropriate for grasping with the real hand. An extensive discussion of the second step of the optimisation strategy follows the next section.

6.3.1 One-/Two-Shot Learning of the Pre-Grasp

In the first step of the optimisation strategy, the pre-grasp is optimised for simultaneity of finger contacts. Presuming the condition that the object is fixated in the simulated world, a *one-shot learning* of the pre-grasp posture leads to simultaneous contacts during grasp closure if all grasping fingers can touch the object. But some fixated benchmark objects are out of reach of the thumb. Under normal conditions, these objects are dragged by other fingers, before the thumb can touch them. For these objects, a *two-shot learning* is applied in which a second learning step optimises the thumb posture, to reduce the elapsed time between the first and the last finger contact.

The **one-shot learning** consists of the following steps:

1. Execute the default pre-grasp of the grasp to be optimised.
2. Apply the target grasp until the first finger has contact with the fixated object.
3. Store the current joint angle values of all fingers.
4. Continue closing the grasp until the next finger touches the object.
5. Subtract the finger's current joint angles from the previously stored ones, and add this difference to its joint angle values of the default pre-grasp.
6. Repeat with step 4 until all grasping fingers have contact or do not move anymore.

When grasping with the Shadow Hand by utilising the power or the all finger precision grasp, it happens that not all grasping fingers touch the fixated object. For the fingers that miss the object, the joint angle values stored in step 3 are taken as the new pre-grasp angles. The resulting pre-grasp posture shows a good approximation to the shape of the target object (see Figures 6.10a and 6.10b). Fingers that did not touch the fixated object can be a strong support in the subsequent grasp closure when the object is un-fixated (compare Figures 6.10c and 6.10d). Even if a finger is not involved in the simulated grasp (for instance, the little finger in Figure 6.10b) its more object

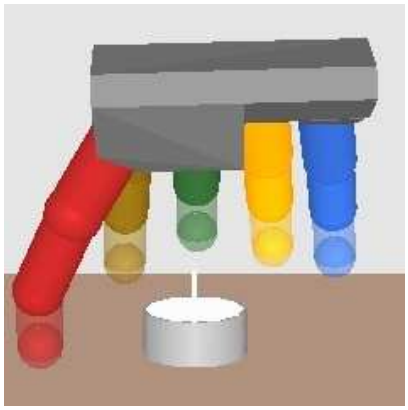

(a) New pre-grasp posture for the *all finger precision grasp* of the tea light.

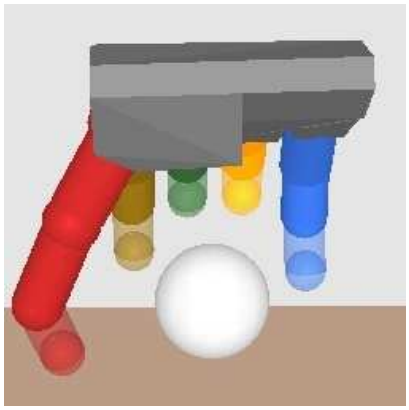

(b) New pre-grasp posture for the *power grasp* of the golf ball.

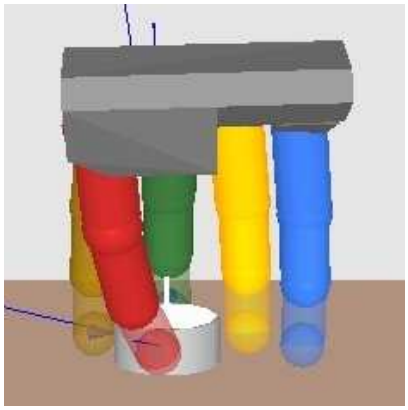

(c) Ring finger closely misses the fixated object.

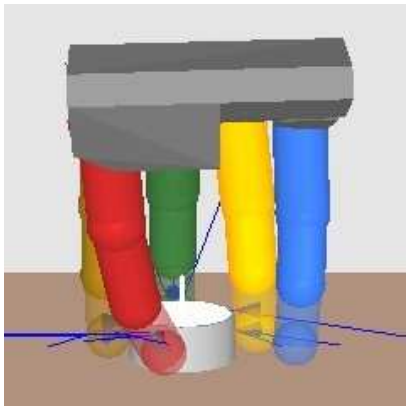

(d) Ring finger supports the subsequent grasp of the un-fixated object.

Figure 6.10: The tea light and the golf ball are two examples of objects that are not touched by all five fingers in the one-shot learning of the pre-grasp.

enclosing pre-grasp posture can be essential for the grasp success of the real hand. The discussion of the grasp evaluation in Section 7.1.2 provides more insights.

If the thumb did not touch the object in the one-shot learning routine, its new pre-grasp posture has to be learned in a second "shot". The other fingers, except for the thumb, touch the object simultaneously when utilising the new pre-grasp angles. The object now has to be un-fixated such that the fingers can drag it towards the thumb.

The **two-shot learning** comprises the following steps:

1. (Perform the one-shot learning first.)
2. Execute the pre-grasp by utilising the new joint angles of the fingers and the default ones of the thumb.
3. Apply the target grasp until the first finger has contact with the un-fixated object.
4. Store the current joint angle values of the thumb.
5. Continue closing the grasp until the thumb touches the object.
6. Subtract the thumb's current joint angles from the previously stored ones, and add this difference to its joint angle values of the default pre-grasp.

The first step of our optimisation strategy results in a *new pre-grasp* for grasping a particular target object. To determine the optimisation progress, the stability value σ of the *new grasp type*, comprising the new pre-grasp and the default target grasp, is evaluated in simulation. Only if the stability value σ of this new grasp is larger than that of the default grasp, the new pre-grasp is used in the second optimisation step.

6.3.2 Optimisation of the Target Grasp

With a hand in default or new pre-grasp posture, the target grasp determines the closing trajectories of the grasping fingers. As already pointed out, trying to optimise all trajectories together is not reasonable. However, this is not necessary. Based on experiences gained during the first evaluation of the standard grasps (compare Section 4.2), and because of the importance of the thumb posture examined in Section 2.2.3, we expect the most benefit from optimising the thumb trajectory.

In most existing robot hands, like the two we have at our disposal, a joint comprising two degrees of freedom located at the thumb's base mimics the saddle joint of the human hand. This joint primarily is responsible for opposing the thumb and thus determines its trajectory during grasp closure. Optimising only the first two thumb joints keeps the optimisation of the target grasp posture as simple as possible. Nevertheless, this is no trivial problem. From experiences of adjusting the two joint angles in the simulator by hand, we know that simply adjusting one joint angle after the other leads to success only occasionally. The two joint angles are correlated, and the two dimensional search space includes a lot of local minima. For optimising the target grasp posture, a search method is needed that finds the global minimum in this search space.

Optimisation and Search Methods

Conventional optimisation and search algorithms can be divided into three main types of methods: *calculus-based*, *enumerative*, and *random*. *Calculus-based* methods depend on the existence of derivatives of the objective function. To formalise the objective function that calculates the grasp stability, we focus a problem of high dimensionality. As described in Section 6.2.4, the grasp stability depends on the quality values of several simulation steps and thus on contact locations, contact wrenches, and the kinematics and dynamics of the simulated system. The formalisation of the objective function and its derivatives is hardly possible, but we can determine its values (grasp stability values σ).

Enumerative search algorithms, like brute-force search, compare objective function values at every point in the search space. This kind of search method usually suffers from the "curse of dimensionality", and is not feasible in large search spaces. When optimising off-line in the two-dimensional search space, this could be a usable method. But our ambition was implementing an optimisation algorithm that is more efficient and extendable to higher dimensions (benchmarking results are presented in Section 6.3.5).

Random search algorithms have achieved increasing popularity as researchers have recognised the shortcomings of calculus-based and enumerative methods. But random searches, in the long run, can be expected to do no better than enumerative methods.

Enumerative and random search methods only explore the search space. More efficient algorithms exploit the current best solution using an iterative improvement technique. The most promising techniques belong to the class of stochastic optimisation methods, like *hill-climbing*, *simulated annealing*, and *evolutionary computation*.

In stochastic gradient descent (*hill-climbing*) a point in the neighbourhood of the current best solution is selected after approximating the gradient of the objective function only by evaluating a single training example. Other forms of this kind of search method use a small number of training examples for approximating the true gradient, but each method only provides local optimum values.

To find the global optimum while being efficient, more sophisticated search methods balance between exploration and exploitation of the search space. *Simulated annealing* [Kirkpatrick et al., 1983] is a technique inspired by annealing in metallurgy. In each iteration step, the previous best solution is replaced by the current point with a probability p, depending on its objective function value and an additional control parameter, "temperature" T. In general, a lower temperature T means a smaller probability p for the current point to be accepted. Beginning with a high temperature T in the first iteration step, this control parameter is lowered in subsequent steps. The search methods alters from exploration to exploitation and may lead to a good solution in a short time. On the other hand, it can be expected that during the execution of the algorithm, the current best solution "jumps" between regions of local minima, before the final search region is chosen when the temperature T is low enough. Thus, only one region is explored accurately. While providing a good solution, the global optimum might be missed.

Like other probabilistic-based algorithms, techniques of *evolutionary computation* do not guarantee the optimal solution. The major difference of evolutionary techniques compared to traditional optimisation and search procedures is that they search from a population of potential solutions, not a single point. Therefore, these methods explore promising regions of local minima more accurately, and a carefully designed algorithm leads to an optimal – or near-optimal – solution while being efficient.

Evolutionary Computation

The field of evolutionary computation comprises different evolutionary-based methods [Fogel, 1994]. Major approaches are: *genetic algorithms*, *evolution strategies*, and *evolutionary programming*. These methods are based on the principle of evolution: "survival of the fittest". Systems build on these approaches maintain a population of individuals, where each individual represents a potential solution to the problem at hand. Random changes based on probabilistic algorithms are imposed to those solutions, and a selection scheme, biased towards fitter individuals, selects which solutions to maintain into the next generation and which to remove from the population.

In analogy to natural evolution, individuals that are selected to be altered are called "parents", and the resulting individuals are called "offspring". The reproduction methods used to create new offspring are termed "genetic operators" and can be divided into "mutation" and "recombination" (also called "crossover"). Differing from the natural model, individuals are often equated to "chromosomes", and these one-chromosome individuals consist of a number of "genes".

Differences between the approaches to evolutionary computation are characterised by the methods employed for selecting new parents and the types of reproduction methods that are imposed on the population to create offspring. Additionally, the representations of the individuals can be different. Each encoding of the individuals involves an appropriate data structure and suitable operators. For example, in binary representations, genes are represented as single bits, and the genetic operators use simple bit-transformations.

Genetic algorithms traditionally used fixed-length binary strings as domain independent representation of the individuals. The beginnings of genetic algorithms can be traced back to the early 1950s. It is the best known technique in the area of evolutionary computation. The work of John

Holland and his colleagues [Holland, 1975] led to genetic algorithms as they are known today. Recombination is used as the primary genetic operator while mutation is secondary.

Evolutionary programming, developed by Fogel et al. [1966], uses representations that are tailored to the problem domain. For example, in real-valued optimisation problems, the individuals are represented as real-valued vectors. For transformation, all individuals are selected and then mutated. Recombination of individuals is not generally performed.

Because *evolution strategies* traditionally were developed for parameter optimisation problems [Rechenberg, 1973], they use real-valued vector representations. Individuals are selected uniformly randomly to be parents. Pairs of parents are recombined to create offspring which are further perturbed via mutation.

These three major approaches in turn have inspired the development of many other approaches in the field of evolutionary computation. For an overview and further references see, for example, Eiben and Smith [2003]. Nowadays the distinction between these approaches is blurring, particularly as far as representation and operators are concerned. Most researchers agree that individuals should be represented in a data structure most suitable to the problem at hand. For example, for real-valued parameter optimisation problems, the parameter values are usually directly stored (in contrast to using a binary encoding), and mutation operators are primarily applied even when the general framework conforms to a classical genetic algorithm.

6.3.3 Design of the Evolutionary Algorithm

We use an evolutionary framework developed in our working group [Henschel, 2006]. In this framework, an evolutionary algorithm can be variably designed. Depending on the design decisions, it can approximate any of the approaches introduced above. Our implementation of an evolutionary algorithm for optimising the thumb posture is described in Section 6.3.4. By means of the single steps of an evolutionary algorithm (see Figure 6.11), design options and aspects that are important for our implementation are explained in this section. For a more detailed description see Henschel [2006].

Each design decision has to consider the balance between exploration and exploitation. When increasing exploitation, the ”selective pressure” is enhanced, which is important for a fast convergence of the algorithm. But premature convergence has to be avoided. More exploratory methods enlarge the ”population diversity” while the selective pressure is reduced. The problem is, to find the optimal design of the evolutionary algorithm for the task at hand. The algorithm has to be efficient, without converging too fast leading to an undesired solution. This balance can be influenced by design choices in each step of the evolutionary process.

The evolution algorithm maintains in each iteration t a population $\boldsymbol{P}(t)$ of a number n of individuals $\boldsymbol{I}$

$$\boldsymbol{P}(t) = \{\boldsymbol{I}_1^t, \ldots, \boldsymbol{I}_n^t\}\,. \tag{6.7}$$

Individuals $\boldsymbol{I}_j^t$ are represented as floating point arrays and consist of a number m of genes G

$$\boldsymbol{I}_j^t = \{G_{j_1}^t, \ldots, G_{j_m}^t\}\,. \tag{6.8}$$

Initialisation

The determination of the population size n is crucial for finding the balance between exploration and exploitation of the evolutionary algorithm. A larger population increases the diversity of

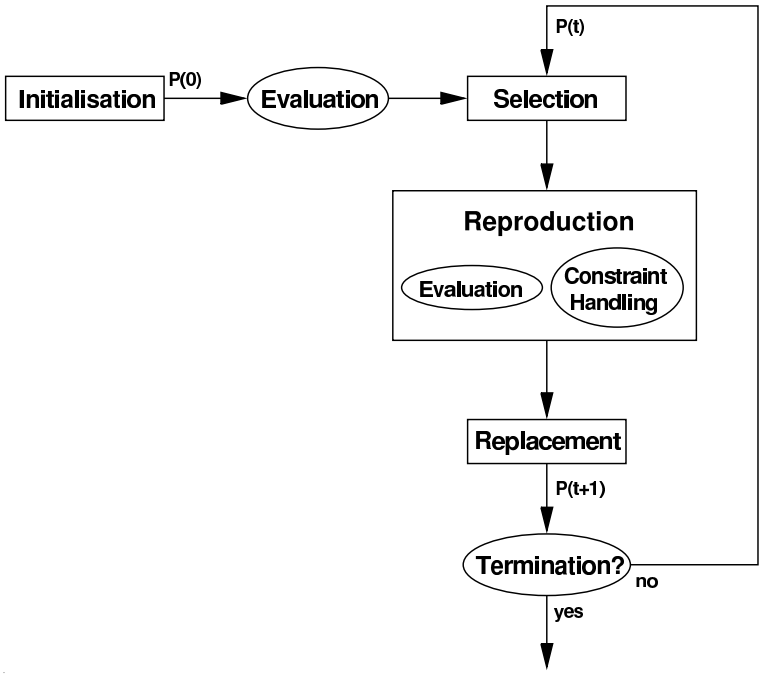

Figure 6.11: Steps of the evolutionary algorithm.

chromosomes, which reduces unidirectional searches and thereby may protect from premature convergence. Heuristics about the number and the regions of potential solutions help to keep the evolutionary algorithm efficient when used for initialising the population ($t = 0$)

$$\boldsymbol{P}(0) = \{\boldsymbol{I}_1^0, \ldots, \boldsymbol{I}_n^0\}\,. \tag{6.9}$$

Methods to initialise the genes of the individuals are:

- *Random with Global Interval*
  Each gene G_k^0 ($k = 1, \ldots, m$) is selected randomly out of an interval i. This interval i has to be specified beforehand and is valid for all genes.
- *Random with Local Intervals.*
  A particular interval i_k ($k = 1, \ldots, m$) has to be specified for each gene G_k. The initial gene G_k^0 is selected randomly out of its interval i_k.
- *Previous Sequence*
  Different sequences can compose the evolution process. For each sequence different design decisions can be made. Subsequent sequences can use the potential solutions of the previous one.
- *Manually*
  Heuristics on good solutions or a promising population diversity can be used by providing an appropriate set of individuals manually.

Evaluation

The evaluation function f is the link between the evolutionary algorithm and the problem to be solved. Each individual $\boldsymbol{I}_j^t$ is evaluated by applying an evaluation function $f(\boldsymbol{I}_j^t)$. It is also called "fitness function" because it returns the "fitness" of the individual, and it corresponds to the

”objective function” in traditional optimisation methods. This fitness function does not have to be known to the evolutionary algorithm. For solving the task at hand, it only needs to get appropriate fitness values of each individual.

Selection

There exist different mechanisms for the selection of which individuals should reproduce. When the population size n is large, it might not be appropriate to select all individuals for reproduction. Selection methods that focus on a smaller part of the population might be required. Methods that are biased towards fitter individuals, enhance the selective pressure of the algorithm.

Four major selection mechanisms are:

- *All Individuals*
  The reproduction mechanism tries to reproduce each individual.
- *Random*
  A number of individuals is selected randomly.
- *Quality Based*
  A number of individuals is selected according to their fitness. Individuals of higher quality are favoured.
- *Tournament Based*
  A number of randomly chosen subsets of individuals is generated. The fittest individual of each subset is selected for reproduction.

Crossover operators require a second or more individuals for reproduction. In which way the crossing partners are chosen, depends on the type of crossover operator applied.

Reproduction

Reproduction is the process of generating offspring by mutating a single individual or crossing different individuals. Different types of both genetic operators, mutation and crossover, exist.

Mutation is asexual reproduction. It generates new genes by chance. Thus, mutation is basically an exploratory reproduction method. Mutating a gene G with a mutation rate r means that it is increased or decreased by r with a 50% chance each. A promising optimisation strategy is, to use large r in the beginning of the evolution process and use smaller r at its end. Large r lead to rapidly finding promising regions of potential solutions of the search space. These findings can be exploited, by lowering the mutation rate r, which aims for local fine tuning.

Major types of the mutation operator are:

- *Basic Mutation*
  This method of reproduction uses fixed mutation rates r_f.
- *Dynamic Mutation*
  Dynamic mutation uses a dynamically changing mutation rate $r(t)$. Beginning with a starting rate r_s it decreases with each iteration step t: $r(t) = r_s/t$.

These two operators are used in our implementation (see Section 6.3.4). Which gene $G^t_{j_k}$ of the selected individual $\boldsymbol{I}^t_j$ is to be mutated, is chosen randomly. Other applicable mutation operators are: boundary mutation, uniform mutation, heavy mutation, random mutation and decay mutation.

Crossover is a recombination method of two or more parents. It provides a mix of genes of the parents to their offspring. The offspring becomes a kind of point between their parents, depending on the crossover method used. By using crossover only, certain regions (in which the global minimum might lay) may not be examined. Even local minima might not be reached because they might lay outside of any centre of the parents. On the other hand, successful genes can be propagated to other members of the population. Thus, crossover is an exploiting reproduction method that increases the selective pressure.

- *Basic Crossover*
  Basic crossover is a two parent crossing mechanism. A certain number of genes is taken from the first individual $\boldsymbol{I}^t_j$ and the rest from the second $\boldsymbol{I}^t_k$. The split point as well as the second individual $\boldsymbol{I}^t_j$ is selected by chance while self crossing is prevented.

- *Average Crossover*
  This method generates its offspring $\boldsymbol{I}^{t+1}_j$ by averaging genes of two parents $\boldsymbol{I}^{t+1}_j = \frac{\boldsymbol{I}^t_j + \boldsymbol{I}^t_k}{2}$. The crossing partner is selected by chance and self crossing is prevented.

Further crossover methods that are not used in our implementation for optimising the thumb posture are heuristic crossover, arithmetical crossover, and poly fixed crossover.

For each operator an **appliance probability** p with $0 < p \leq 1$ has to be specified. The chance of an operator for being applied depends on p and the selected method for offspring insertion (explained below in Section "Replacement").

Constraint Handling

Constraints, like bounds of gene values, can be provided for each gene. This is useful for preventing the evaluation function from being passed infeasible values. Additionally, the search space can be restrained by excluding unpromising regions, which speeds the optimisation procedure. A constraint handling mechanism decides how to handle constraint violation.

- *Death Penalty*
  This mechanism rejects an infeasible individual $\boldsymbol{I}^{t+1}_j$ if a gene $G^{t+1}_{j_k}$ is generated that violate the constraints (for example, its value lays outside of a given interval i_k). The Death Penalty method is useful when only simple constraints are specified, for example, if the solution space is limited by single gene intervals.

When several intervals are specified for each gene, the solution space is more complex, and more sophisticated constraint handling methods may be required for finding the region of the global optimum. These methods accept different amounts of constraint violations. They provide a penalty component which reduces the quality of infeasible solutions depending on the amount of their violation. Applicable constraint handling mechanisms are: Customised Death Penalty, Feasible Infeasible Penalty, Linear Best Penalty, Dynamic Penalty, Adapted Penalty, and Better Best Penalty [Henschel, 2006].

Replacement

Whether an offspring is inserted into the population depends on the replacement technique used:

- *In-Place Offspring*
  Each offspring $\boldsymbol{I}_j^{t+1}$ directly replaces its parent $\boldsymbol{I}_j^t$ according to its quality. It has to be pre-defined how many attempts a a parent can be reproduced until an appropriate quality offspring is reached. It can be selected that degeneracy is allowed, which means that a lower quality of offspring compared to that of its parent is accepted. In each trial (1 to a) one of the specified operators is applied while operators with higher probability value p have more chance to be chosen.

- *Best Individuals*
  New offspring is inserted into the population only if it performs better than the individual with the lowest fitness value. Before creating the next generation $\boldsymbol{P}(t+1)$ by inserting offspring, reproduction is performed on all individuals selected for this purpose. On these individuals each specified operator is applied with the probability p. The best n individuals of the set of the previous generation $\boldsymbol{P}$, together with all new offspring, form the next generation $\boldsymbol{P}(t+1)$.

Best Individuals is a highly exploitative method. Choosing only the fittest individuals for the next generation $\boldsymbol{P}(t+1)$ enforces a fast convergence of the algorithm. In-Place Offspring allows a parallel development of different potential solutions and has a lower selective pressure. Other, but not used, replacement techniques are:

- *Fixed Insertion*
  Offspring always is inserted, a randomly selected individual is removed from the population.

- *Total Replacement*
  The reproduction process generates n offspring, which form the next generation $\boldsymbol{P}(t+1)$.

Termination

When the next generation $\boldsymbol{P}(t+1)$ is created, a new iteration step $t+1$ starting with the selection process can be executed. After which iteration step the evolutionary algorithm terminates depends on the selected termination condition:

- *Threshold Condition*
  The evolutionary algorithm terminates once the average quality increase of the complete population $\boldsymbol{P}(t+1)$ over a given number of generations g is lower than a chosen threshold value $q > 0$ with

$$\sum_{j=1}^{n} Q(\boldsymbol{I}_j^{t+1}) - Q(\boldsymbol{I}_j^{t+1-g}) < q, \tag{6.10}$$

  where Q returns the individual's quality, which previously was determined by the fitness function $f(\boldsymbol{I}_j^t)$.

- *Generation Condition*
  The evolutionary algorithm terminates after a given number of iterations t_{max}.

6.3.4 Implementation for Thumb Angle Optimisation

The design of an evolutionary system depends on the kind of problem to be solved and the experience of the user. An appropriate balance of exploration and exploitation is required for good performance. Especially choosing mutation and crossover operators may or may not be appropriate for the task at hand. Each operator plays a different role in the search process. A priori, it is difficult to specify the relative importance of each operator. The choice of appropriate operators is rather a heuristic process, and even experienced users test different operators and other design possibilities before determining the final configuration of an evolutionary system [Spears, 1992].

Our evolutionary algorithm for optimisation of the first two thumb angles is based on the design considerations discussed in Section 6.3.3. Before we tried out several combinations of operators with different parameterisations, we determined all other design decisions and how many sequences our evolutionary algorithm comprises.

Evolution Sequences

The evolutionary algorithm can consist of several sequences. Each sequence constitutes an evolutionary sub-process. Two sequences can represent totally different evolutionary algorithms in that for each step of the algorithm different design decisions can be made. A constructive sequence improves the potential solutions of the terminated one.

In our implementation for optimisation of the first two thumb angles, we use two sequences with different purposes. The first sequence has to explore the solution space while maintaining the initial population diversity. The purpose of the second sequence is to exploit the potential solutions of the first sequence while enlarging the selective pressure and exploring the most promising regions. The desired behaviour of the evolutionary algorithm is achieved with following design decisions.

Design Decisions

There are several design decisions of different steps of the evolutionary algorithm that apply for both sequences. We choose a **population size** of $n = 14$ individuals. Each individual $\boldsymbol{I}_j$ consists of $m = 2$ genes (the first two joint angles of the thumb).

Selection: *All Individuals* are selected for reproduction.

Evaluation: For evaluating the fitness of an individual, we execute the grasp to be optimised in simulation while replacing the first two thumb joint angles of the target grasp with the gene values of the individual. The stability measure described in Section 6.2.4 provides the fitness value of the individual to be evaluated.

Constraint Handling: We use single intervals i_k for both genes G_k ($k = 1, 2$) to limit the search space to promising regions. These limits depend on the used grasp type and are presented in Table A.3.

Death Penalty is applied as the constraint handling mechanism. Thus, constraint violation is not allowed.

Termination: Because we optimise off-line and search for the optimum solution, we use *Threshold Condition* for termination with a very low threshold of $q = 0.0001$. The number of generations to be sampled over is $g = 3$.

The design decisions in steps: Initialisation, Reproduction, and Replacement differ between the two sequences:

1. Sequence

Initialisation: The method used to initialise the $n = 14$ individuals $\boldsymbol{I}_j^0$ is *Manually*. The individuals are distributed over the two-dimensional search space in accordance to the intervals i_k presented in Table A.3.

Reproduction: We use four differently parameterised mutation operators to explore the search space for potential solutions:

- *Basic Mutation* with mutation rate $r_f = 1$, appliance probability $p = 0.15$;
- *Basic Mutation* with mutation rate $r_f = 0.5$, appliance probability $p = 0.15$;
- *Basic Mutation* with mutation rate $r_f = 0.25$, appliance probability $p = 0.15$;
- *Basic Mutation* with mutation rate $r_f = 0.1$, appliance probability $p = 0.15$.

Replacement: By utilising the *In-Place Offspring* technique in the first sequence, the next generation $\boldsymbol{P}(t+1)$ maintains individuals that remain in similar regions to those of the previous generation $\boldsymbol{P}(t)$. Together with the uniform distribution of the initial population, an accurate exploration of different regions of potential solutions is possible. Maximal $a = 3$ attempts of achieving a larger fitness value by applying a single operator on each individual $\boldsymbol{I}_j^t$ are executed. Degeneration is not allowed.

2. Sequence

Initialisation: After termination of the first sequence, its resulting population is employed as the initial population $\boldsymbol{P}(0)$ of the second sequence by utilising the *Previous Sequence* method.

Reproduction: Besides four basic mutation operators with lower mutation rates than applied in the first sequence, one additional dynamic mutation and two crossover operators aim for exploiting the potential solutions and for fine tuning:

- *Basic Mutation* with mutation rate $r_f = 0.08$, appliance probability $p = 0.15$;
- *Basic Mutation* with mutation rate $r_f = 0.05$, appliance probability $p = 0.15$;
- *Basic Mutation* with mutation rate $r_f = 0.025$, appliance probability $p = 0.15$;
- *Basic Mutation* with mutation rate $r_f = 0.01$, appliance probability $p = 0.15$;
- *Dynamic Mutation* with starting rate $r_s = 1$, appliance probability $p = 0.15$;
- *Basic Crossover* with appliance probability $p = 0.15$;
- *Dynamic Crossover* with appliance probability $p = 0.15$.

Replacement: Only the n fittest individuals are chosen for the next generation $\boldsymbol{P}(t+1)$, by utilising the *Best Individuals* technique. This enlarges the selective pressure, and after a few iteration steps it is expected that all n individuals concentrate on the most promising region. The local fine tuning results in the desired optimal – or at least near optimal – solution.

6.3.5 Benchmarking the Evolutionary Algorithm

We decided to utilise an evolutionary algorithm for optimising grasp postures because this is an efficient search method even when facing higher dimensions [Michalewicz, 1996]. The first two joint angles of the thumb provide the most potential when optimising the target grasp because they determine the thumb opposition posture. But other joint angles and especially the parameters that determine the position and orientation of the hand can be reasonable parameters to be optimised. The evolutionary algorithm applied can easily be extended to consider individuals of higher dimensions.

Besides the advantage of being extendable, the efficiency of the implemented evolutionary algorithm, i.e. how fast it converges towards an optimal solution, is an important characteristic. Because we use the evolutionary algorithm for optimising two-dimensional individuals only, it is questionable whether it outperforms conventional search algorithms, like enumerative and random methods, which are most advantageous when applied in low dimensional search spaces. To answer this question, we performed a benchmarking of our evolutionary algorithm against a brute-force search method.

The evolutionary algorithm implemented requires a fitness value after each application of an operator generating a new offspring. The evaluation of an individual, by simulating the grasp and calculating the grasp stability value σ, is the time consuming factor of the optimisation process.[2] Thus, an enumerative search algorithm that evaluates a similar number of points (equals individuals) requires a similar amount of time.

objects	grasp type	EA		brute force search			
		indiv.	σ	points	stability	points	σ
TUM Hand							
toy propeller	3F spec t_5	1234	0.808905	1260	0.423482	5130	0.654956
remote control	power t_3	1096	0.311009	1147	0.298496	5070	0.305300
light bulb	power t_3	1012	0.285743	1015	0.271614	5070	0.263055
folding rule	2F prec t_2	967	0.125237	999	0.112643	5063	0.117694
pencil	2F prec t_2	974	0.326044	999	0.247617	5063	0.214468
Shadow Hand							
tissue pack	power t_3	1121	0.240717	1147	0.230789	5082	0.237272
board marker	2F prec t_2	1221	0.056605	1271	0.043741	5084	0.046284
chocolate bar	AF prec t_1	1070	0.274517	1089	0.266396	5041	0.261528
eraser	2F prec t_2	1153	0.24114	1200	0.136821	5084	0.140026
bunch of keys	AF prec t_1	1268	0.260685	1296	0.256091	5041	0.262884

Table 6.1: Benchmarking the implemented evolutionary algorithm (EA) against brute force search. When providing the brute force search with a number of points similar to the number of individuals (indiv.) evaluated for the EA, the EA outperforms the enumerative search in each case (red stability values σ). Only when grasping the bunch of keys with the Shadow-Hand while evaluating more than 5000 points, the brute force search leads to a higher stability value σ (when using 4096 points in this case, $\sigma = 0.255958$).

We determined the efficiency of the evolutionary algorithm by comparing the fitness value of the best individual with the fitness value of the best point of a brute-force search. The number of points evaluated in this enumerative search algorithm was chosen to be uniformly distributable over the two-dimensional search space and to be larger than the number of offspring evaluated

[2] Depending on the used grasp type and the target object, evaluation of one individual takes between 0.4 seconds and 1.4 seconds on a 1.8 GHz Pentium 4 CPU.

before the best individual was found with the evolutionary algorithm. The search was limited by the bounds of the genes, similar as when using the evolutionary algorithm. We chose ten different objects that were grasped with different grasp types. Each robot hand grasped five of them. In all ten cases the evolutionary algorithm outperforms the brute force search, because when evaluating a similar number of points, the evolutionary algorithm found an individual with a larger stability value (see Table 6.1).

To determine the efficiency of the evolutionary algorithm, we enlarged the number of points of the enumerative algorithm in steps of thousands. When providing the brute force search with more than 5000 points, for one grasp (bunch of keys with the Shadow Hand) a point was found with a stability value σ larger than the best result of the evolutionary algorithm. On the other hand, in nine of ten cases the brute force search leads to a worse result than the evolutionary algorithm, although about five times as much points were evaluated requiring five times as much time. In other words, in these cases the evolutionary algorithm was five times as efficient as the brute force search while leading to a better result.

7 Evaluation

Based on the benchmark system proposed in Section 4.1, we performed several experiments in both robot hand setups to evaluate the optimisation strategy proposed in Section 6.3.

These experiments are described and analysed in Section 7.1. The meaning of the experimental results on the different steps of the optimisation strategy is evaluated in Section 7.2. The results of one additional experiment lead to a comparison of the grasping capabilities of the robot hands we used as described in Section 7.3.

7.1 Experiments

In the experiments we performed, the success rates (0 to ten, out of ten grasp trials) for grasping each object were determined. The first experiment was conducted for evaluating the grasp success of the standard grasps. After optimising the pre-grasps of the standard grasps in simulation, the results (*new pre-grasp*) were analysed in terms of their grasp stabilities σ and were used to conduct the second experiment with the real hand. If the stability of a new pre-grasp was larger than that of the standard grasp, the new pre-grasp was used in the second optimisation step. Otherwise, the standard pre-grasp was utilised.[1] The chosen pre-grasp was applied in the third experiment with the real hand together with the optimised target grasp (*new target grasp*) which was the result of the second optimisation step. Detailed listings of the optimised values of the joint angle configurations are presented in Section A.2.

The evolutionary algorithm applied in the second optimisation step requires a fitness value (stability σ) suitable for finding promising search regions (see Section 6.3.4). To this end, we enlarge the frictional coefficients μ and ν (see Section 6.1.2) if the stability $\sigma < 1e^{-6}$ after the first grasp optimisation step. In the search space of the evolutionary algorithm, this leads to more grasps with quality values α larger than zero (see Section 6.2) and thus to more grasps with stability σ larger than zero. This enhances the chance for finding promising search regions and the optimal target grasp.

The grasp type t to be chosen for grasping a particular object with the Shadow Hand was determined to be the same type that was identified by the first evaluation of the TUM Hand standard grasps as described in Section 4.2. This provides the opportunity to compare grasping of a particular benchmark object between both robot hands by applying grasps of the same grasp type t.

Tables 7.1 and 7.2 show the results of the grasp experiments for all objects and each of the three grasps, *standard grasp*, *new pre-grasp*, and *new target grasp*. In the case of the TUM Hand, the numbers of successful trials when applying the standard grasps were determined in the first experiment described in Section 4.2 and are redisplayed in the column "standard grasp success" of Table 7.1. The next two sections provide a discussion of the simulated grasps and the experiments performed with the real hands and further explain the results given in Tables 7.1 and 7.2.

[1] One object, the cup, is an exception. For explanation, see the detailed discussion of grasping this object.

7.1.1 Simulated and Real Grasps of the TUM Hand

no.	object	grasp type	standard grasp		new pre-grasp		new target grasp	
			stability σ	success	stability σ	success	stability σ	success
1	adhesive tape	t_3	0.004892	10	0.112395	10	0.242452	10
2	toy propeller	t_5	0.126676	10	0.178193	10	0.808905	10
3	toy cube	t_4	0.093513	10	0.093513	(10)	–	–
4	can	t_3	0.109237	10	0.114548	10	0.186490	10
5	tissue pack	t_3	0.386662	10	0.460620	10	0.486769	10
6	tennis ball	t_3	0.065507	10	0.208380	10	0.240189	10
7	paper ball	t_3	0.379838	9	0.410645	10	0.522070	10
8	sharpener	t_1	0	8	0.191164	10	0.226848	10
9	remote control	t_3	0.129637	8	0.196658	10	0.311009	10
10	cup	t_3	0.287787	9	0.353118	–	0.356003	10
11	board marker	t_2	0.031801	7	0.000518	0	0.131739	10
12	tea light	t_1	0.125607	6	0.169162	10	0.344191	10
13	golf ball	t_3	0.198839	7	0.268729	10	0.364540	10
14	matchbox	t_1	0.157896	7	0.191704	10	0.304353	9
15	light bulb	t_3	0.097098	6	0.110267	9	0.285743	10
16	chocolate bar	t_1	0.154287	5	0.327594	10	0.365321	10
17	folding rule	t_2	0.094767	4	0.001160	2	0.125237	10
18	voltage tester	t_2	$< 1e^{-6}$	3	$< 1e^{-6}$	2	0.149888	9
19	eraser	t_2	0.069555	4	0.061848	6	0.353659	10
20	bunch of keys	t_1	0.149508	0	0.053156	0	0.163833	0
21	pencil	t_2	0	0	0	0	0.326044	0

Table 7.1: Grasp stability σ of simulated grasps and the number of successful grasp trials (0 to 10) out of 10 grasp attempts with the real TUM Hand before optimisation (standard grasp), after the first optimisation step (new pre-grasp), and after the second optimisation step (new target grasp). Red values indicate *new pre-grasps* that are used in the second optimisation step.

Evaluation of the experiments on grasping the benchmark objects with the *standard grasp*, after the first optimisation step (*new pre-grasp*), and after the second optimisation step (*new target grasp*):

1. adhesive tape:

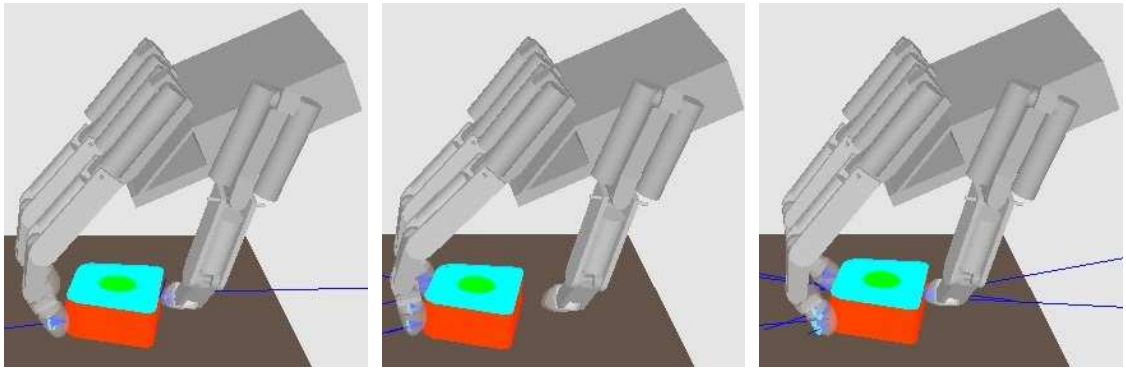

(a) *standard grasp:* thumb pushes the object laterally.

(b) *new pre-grasp:* simultaneous touch of index and middle finger.

(c) *new target grasp:* stable grasp with optimal thumb posture.

Figure 7.1: TUM Hand grasping the adhesive tape in simulation.

- *standard grasp:* When executing the standard power grasp g_3 in simulation, the index finger touches the adhesive tape first and drags it towards the thumb. The second contact is made by the moving thumb which pushes the object laterally because the middle finger cannot provide an opposing force at this point in time (see Figure 7.1a). When finally all three fingers touch the adhesive tape, the contact point of the thumb is suboptimal and leads to a low grasp stability σ. Nevertheless, the evaluation with the real TUM Hand revealed that the standard power grasp was sufficient to grasp the object in each of the ten trials successfully.
- *new pre-grasp:* The first grasp optimisation step generates a *new pre-grasp* which prevents the adhesive tape from being pushed laterally in simulation. Because the index and the middle finger touch the object and drag it towards the thumb simultaneously (see Figure 7.1b), the middle finger now can provide the oppositional force required when the thumb reaches the object. The resulting stability value is much greater than for the *standard grasp*. Again, each grasp trial was successful.
- *new target grasp:* With the second optimisation step, the grasp stability σ was enhanced largely. Since the thumb provides the opposing forces to the other fingers, the contact point on the object is essential for stable grasping (see Figure 7.1c). However, the experiment with the real TUM hand did not show differences to the evaluation of the *new pre-grasp*, and no grasp trial was incomplete again.

2. toy propeller:

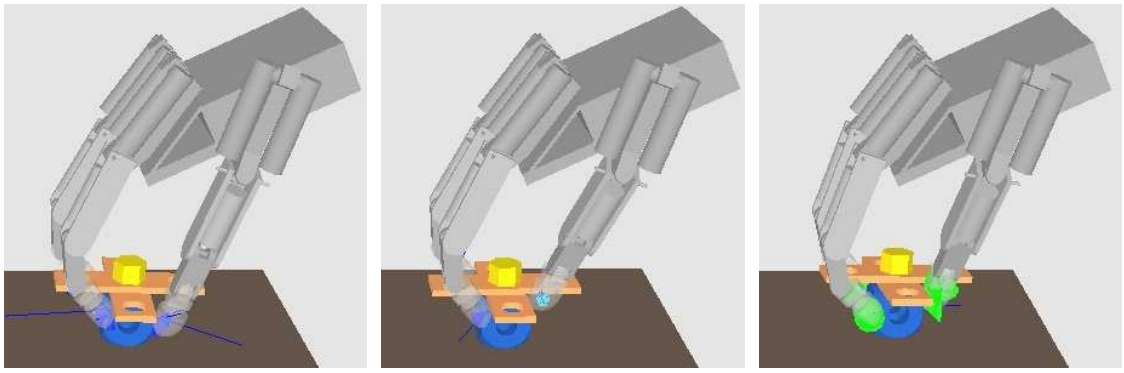

(a) *standard grasp:* thumb touches the toy propeller at the side of a blade.

(b) *new pre-grasp:* simultaneous contacts between the blades.

(c) *new target grasp:* contacts at the cube with wide static friction cones.

Figure 7.2: TUM Hand grasping the toy propeller in simulation.

- *standard grasp:* The simulation of the standard grasp shows that the thumb touches the toy propeller at the side of a blade (see Figure 7.2a). When grasping with the real hand, there is the risk that the thumb misses the blade in the grasp closure phase due to uncertainties in location and orientation of the object. However, during the experiment, the thumb touched the blade at its side in each trial and all grasps were successful.
- *new pre-grasp:* By utilising the *new pre-grasp*, simultaneous contacts are achieved in simulation with each finger touching the propeller between two blades (see Figure 7.2b). This leads to a greater stability σ, and the risk that the thumb may miss the blade is reduced when grasping with the real hand.
- *new target grasp:* The stability value of the grasp was increased largely after optimising the thumb joint angle values. Due to the fact that the object is dragged towards the

wrist before all fingers touch it, the contact points are located at the propellers cube nearer to the object's centre of mass. Furthermore, a static posture is reached in an early simulation step shown by static frictions cones in Figure 7.2c.

3. toy cube:

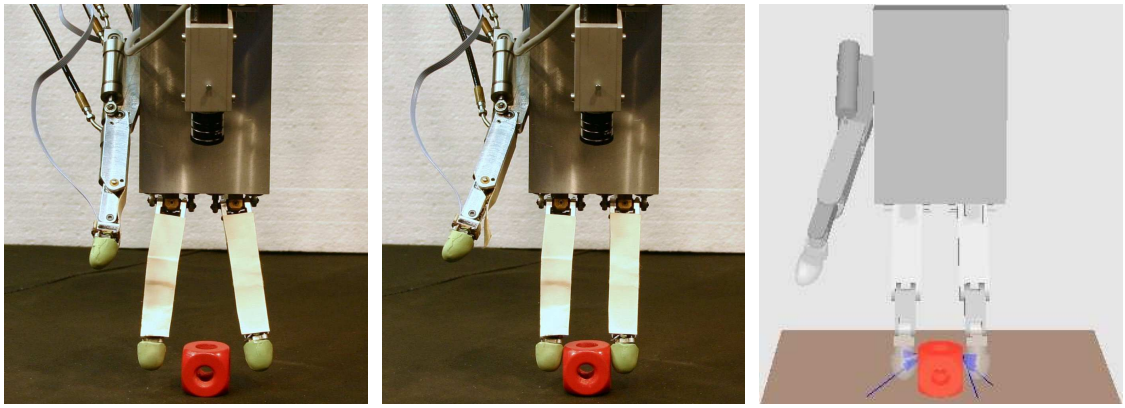

(a) *standard grasp:* the pre-grasp posture is optimal already.

(b) *standard grasp:* a suitable grip posture is achieved.

(c) *standard grasp:* simultaneous contacts of the two grasping fingers.

Figure 7.3: TUM Hand grasping the toy cube in reality and in simulation.

- *standard grasp:* The two finger pinch grasp for grasping the toy cube is the only grasp that cannot be improved by any of the two optimisation steps. Neither this was necessary since the evaluation with the real hand showed that the standard pre-grasp (see Figure 7.3a) and the standard target grasp lead to a suitable grip posture (see Figure 7.3b) for grasping the toy cube successful in each trial.
- *new pre-grasp:* The optimised pre-grasp provides the same joint angles as the standard pre-grasp. This is due to the fact that the two fingers touch the toy cube at the same point in time (see Figure 7.3c) when applying the standard grasp in simulation. A stability evaluation of a grasp with the same preconditions results in the same stability value σ. Thus, a subsequent grasp evaluation was omitted and the brackets around the number of successful grasps in Table 7.2 denote that the success rate of the standard grasp is assumed.
- *new target grasp:* Since the thumb is not involved in the two finger pinch grasp of the TUM Hand, this optimisation step cannot be applied.

4. can:

- *new target grasp:* The closure trajectory of the thumb is more directed towards the object than in the case of the standard power grasp. This is due to a difference in the thumb joint angle θ_1 of more than 20 degrees (compare Tables A.1 and A.4). The result is a greater stability value and the evaluation with the real hand showed full grasp success.

5. tissue pack:

- *new target grasp:* The differences between the standard grasp and the optimised grasps are small in the case of the tissue pack. Every grasp trial was successful

6. tennis ball:

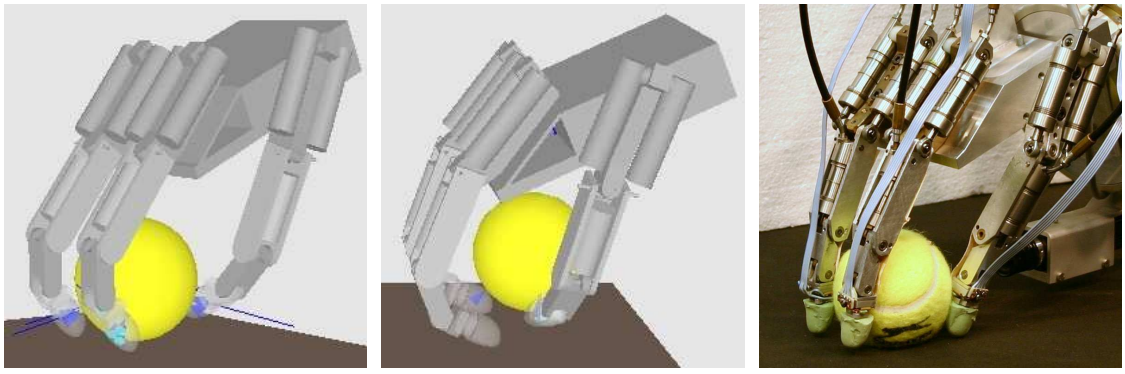

(a) *standard grasp:* object is grasped... (b) *standard grasp:* ...and it is lost afterwards. (c) *new target grasp:* real fingers stop motion when contact is detected.

Figure 7.4: TUM Hand grasping the tennis ball in simulation and in reality.

- *standard grasp:* When the tennis ball is touched by all three finger in simulation (see Figure 7.4a), a grasp suitable for lifting the object is established ($\alpha > 0$). But when the fingers move on towards the target grasp posture, the object is squeezed out of the hand after a few simulation steps (see Figure 7.4b). In contrast, the fingers of the real hand are stopped when contacts are detected (see Figure 7.4c), and in no grasp trial of the experiment the object was lost.
- *new pre-grasp:* With the *new pre-grasp*, the fingers still loose the tennis ball but in a later simulation step. This enhances the stability value greatly, while no essential differences were observed when evaluating with the real hand.
- *new target grasp:* The closure trajectory of the thumb is more directed towards the object which is not lost in simulation anymore. Additionally, the fingers contact the object more simultaneously, and this inhibits the roll movement of the tennis ball. This effect was observed when grasping with the real hand too.

7. paper ball:

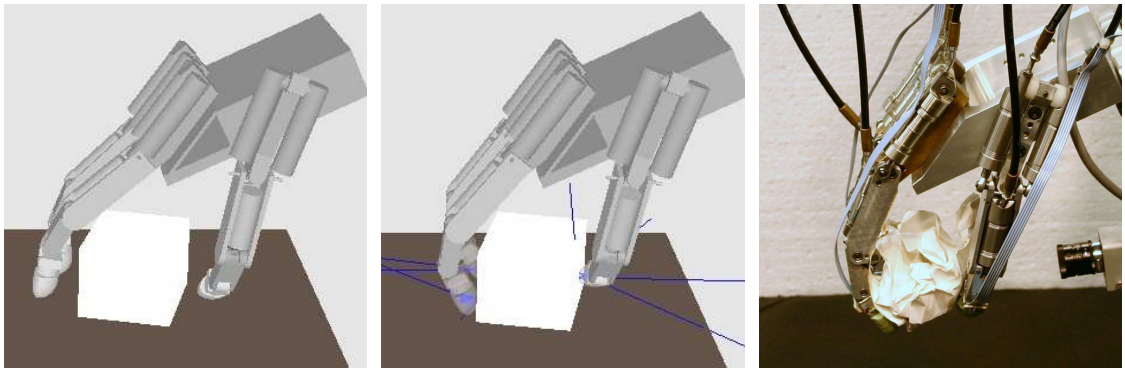

(a) *new pre-grasp:* optimised pre-grasp posture... (b) *new pre-grasp:* ...leads to a suitable grip posture. (c) *new target grasp:* stable grasp with the real hand.

Figure 7.5: TUM Hand grasping the paper ball in simulation and in reality.

- *new pre-grasp:* The cuboid model approximating the paper ball in simulation (see Section 6.1.2) differs largely in its properties from the folded, cleft, compliant sheet of paper. However, with the *new pre-grasp*, a success rate of ten out of ten trials was achievable. The reason is that the further flexion of the thumb in the optimised pre-grasp posture (see Figure 7.5a), in contrast to the standard pre-grasp posture, leads to an optimal thumb opposition posture and a higher stability σ in the subsequent grasp closure (see Figure 7.5b).
- *new target grasp:* After the second optimisation step, no uncertainties were observed when grasping with the real TUM Hand. The optimised thumb angles only differ slightly, but the new target grasp seems to be more stable when grasping with the real hand (see Figure 7.5c), and no ”near-losts” were observed like before.

8. sharpener:

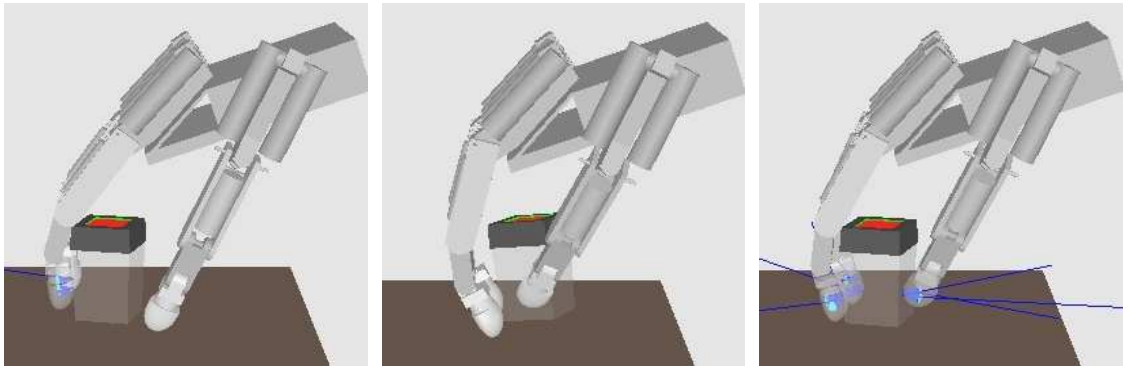

(a) *standard grasp:* index finger touches object at an edge. (b) *standard grasp:* object is lost afterwards. (c) *new pre-grasp:* optimised postures lead to a successful grasp.

Figure 7.6: TUM Hand grasping the sharpener in simulation.

- *standard grasp:* When executing the standard grasp in simulation, first contact is made by the index finger at an edge of the sharpener (see Figure 7.6a). At that time, no other finger touches the object. It is pushed and rotated in a way that the thumb touches the sharpener subsequently at a side on which no forces opposing those of the other fingers can be exerted (see Figure 7.6b). Thus, the sharpener cannot be grasped in simulation and $\sigma = 0$. When grasping with the real hand, the index finger only by chance touches the sharpener at an edge, and the effects described were only observed in two trials. In the other eight grasp trials, the thumb touched the object at a side providing the opportunity to exert opposing forces to the other two fingers.
- *new pre-grasp:* The *new pre-grasp* leads to simultaneous contacts of the index finger and the middle finger. This prevents the sharpener from being rotated. The thumb is more flexed and touches the object at a side appropriate for providing opposing forces to the other fingers. With this optimised pre-grasp, the sharpener can be grasped in simulation (see Figure 7.6c), and with the real hand, each of ten trials was successful.

9. remote control:

- *new pre-grasp:* The *new pre-grasp* leads to a larger grasp stability σ, and the grasp success with the real hand was enhanced from eight to ten out of ten grasp trials (see Figure 7.7a). Although being no topic of this experiment, the previously observed lack of robustness against rotation of the hand after lifting it (see Table 4.3) was eliminated additionally (see Figure 7.7b).

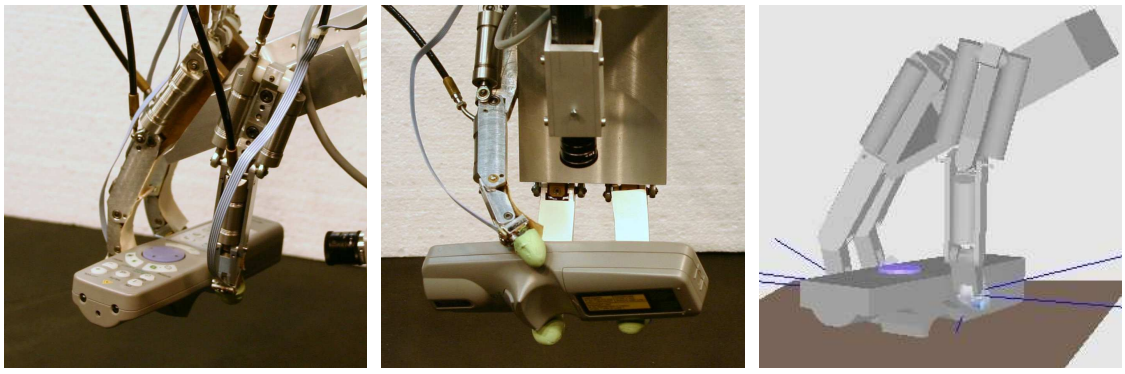

(a) *new pre-grasp:* grasp with large stability. (b) *new pre-grasp:* robustness against rotation is enhanced. (c) *new target grasp:* thumb touches object near its CM.

Figure 7.7: TUM Hand grasping the remote control in reality and in simulation.

- *new target grasp:* After the second optimisation step, the thumb closure trajectory is more directed towards the object's centre of mass (CM) than before (see Figure 7.7c). This results in a larger stability σ and in certain grasp trials with the real hand.

10. cup:

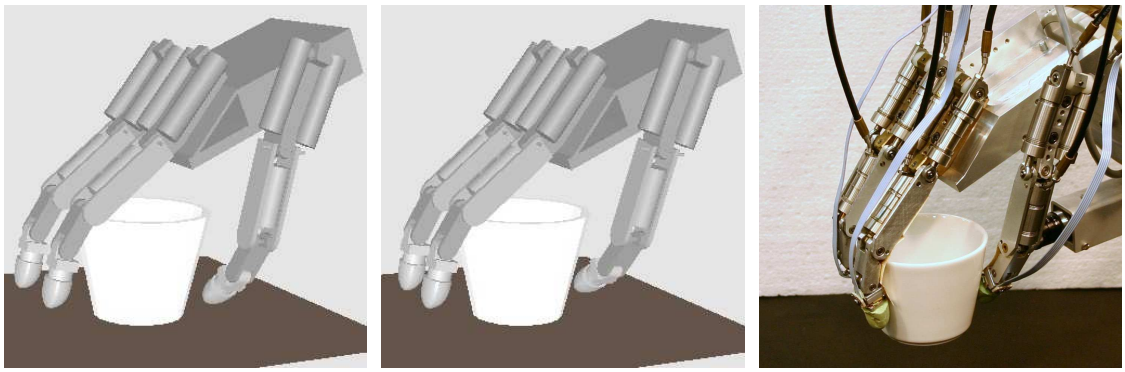

(a) *standard grasp:* pre-grasp posture with a hand opening being large enough. (b) *new pre-grasp:* thumb is flexed too far. (c) *new target grasp:* successful grasp trial.

Figure 7.8: TUM Hand grasping the cup in simulation and in reality.

- *new pre-grasp:* The optimised pre-grasp results in a grasp with larger stability σ than that of the standard grasp (see Figure 7.8a), but it cannot be applied with the real TUM Hand. The reason is that the real hand is in pre-grasp posture during the placing phase, and the thumb sticks into the cup because it is flexed too far in the case of the *new pre-grasp*. In simulation, the grasp evaluation starts with a hand being in pre-grasp posture and in grasp position (see Figure 7.8b). If this opportunity would exist for the real hand too, the optimised pre-grasp could enclose the cup and would result in a more stable grasp. But since the pre-grasp has to be applied before the grasp position is approached, the optimised pre-grasp was not evaluated (indicated with a dash in Table 7.1).
- *new target grasp:* Although the standard pre-grasp is utilised, the new thumb joint angle values lead to a grasp stability σ which is even larger than that of the *new pre-grasp*. As a result, the success rate was enhanced to ten out of ten grasp trials (see Figure 7.8c).

11. board marker:

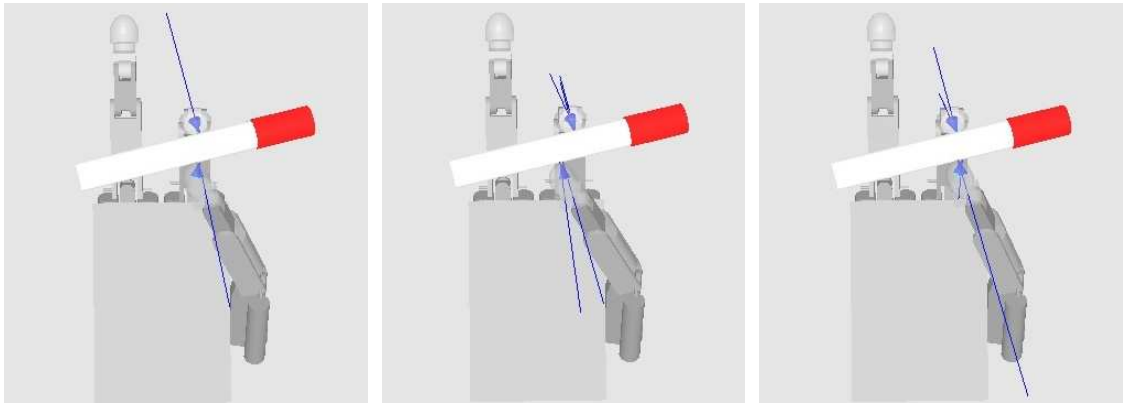

(a) *standard grasp:* quite good standard grasp. (b) *new pre-grasp:* relocated thumb contacts lead to low stability σ. (c) *new target grasp:* optimised opposing forces.

Figure 7.9: TUM Hand grasping the board marker in simulation.

- *standard grasp:* The standard two finger precision grasp is quite suitable for grasping the board marker (see Figure 7.9a). With the real hand seven out of ten grasp trials were successful.
- *new pre-grasp:* The stability σ of the optimised pre-grasp is lower than in the case of the standard pre-grasp. The significant difference between the two pre-grasps is that the optimised one does imply a farther flexion of the thumb. Because the pre-grasp posture of the thumb is rather estimated than optimised when performing a two-shot learning (see Section 6.3.1), this optimisation step does not provide an optimal result in the case of grasping the board marker (see Figure 7.9b). With the real hand, in none of the grasp trials the object was grasped successful, while in most cases the object rotated between the fingers and was lost during the lift-off phase.
- *new target grasp:* The optimisation of the thumb target grasp posture leads to optimal opposing contacts (see Figure 7.9c), and the success rate was enhanced to the optimum.

12. tea light:

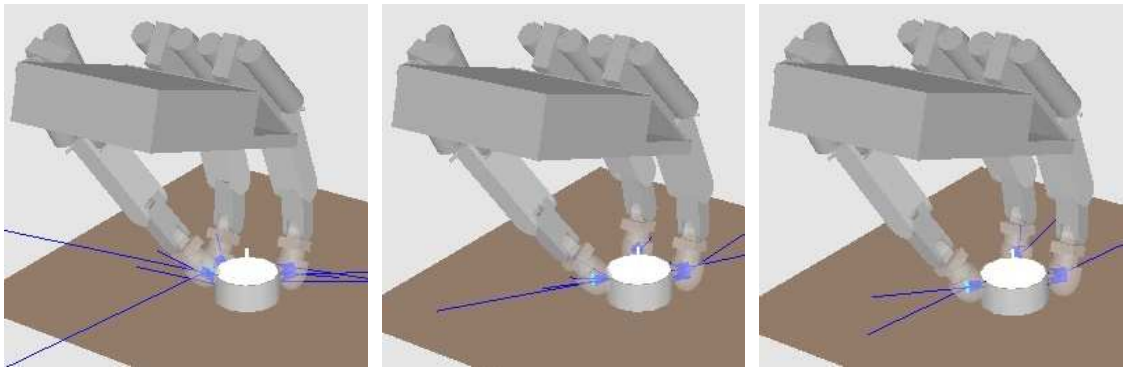

(a) *standard grasp:* object is grasped but lost afterwards. (b) *new pre-grasp:* object is better enclosed by thumb and middle finger. (c) *new target grasp:* thumb exerts optimal opposing forces.

Figure 7.10: TUM Hand grasping the tea light in simulation.

- *standard grasp:* In simulation, the tea light is grasped at first (see Figure 7.10a), but because the thumb cannot exert large enough opposing forces to the other fingers, the object is pushed towards the ulnar side of the hand and lost afterwards. The same effect was observed in the incomplete grasp trials executed with the real hand.
- *new pre-grasp:* Because in the *new pre-grasp* posture the thumb and the middle finger are more flexed than in the standard pre-grasp posture, in the subsequent grip posture, the forces are more synchronically distributed over the side surface of this cylindrical object (see Figure 7.10b). Therefore, in simulation and in reality the tea light is not pushed out of the hand anymore.
- *new target grasp:* As shown in Figure 7.10c, the thumb touches the object at a lower point which facilitates optimal opposing forces to the other finger contacts. This increases the stability σ largely. The grasp trials with the real hand showed no uncertainty, and again full success was achieved.

13. golf ball:

- *new pre-grasp:* Similar to the *new pre-grasp* of the tea light, the middle finger is more flexed compared to the standard pre-grasp posture. This prevents the golf ball from rolling to the ulnar side and out of the hand when grasping with the real hand. The results are an enhancement from formerly seven to ten successful grasp trials and a larger stability σ.
- *new target grasp:* When utilising the new thumb joint angle values, the thumb is more flexed in its target grasp posture, and its closure trajectory is more directed towards the object's centre of mass. This leads to a greater stability value in simulation, whereas again full grasp success was reached when evaluating with the real hand.

14. matchbox:

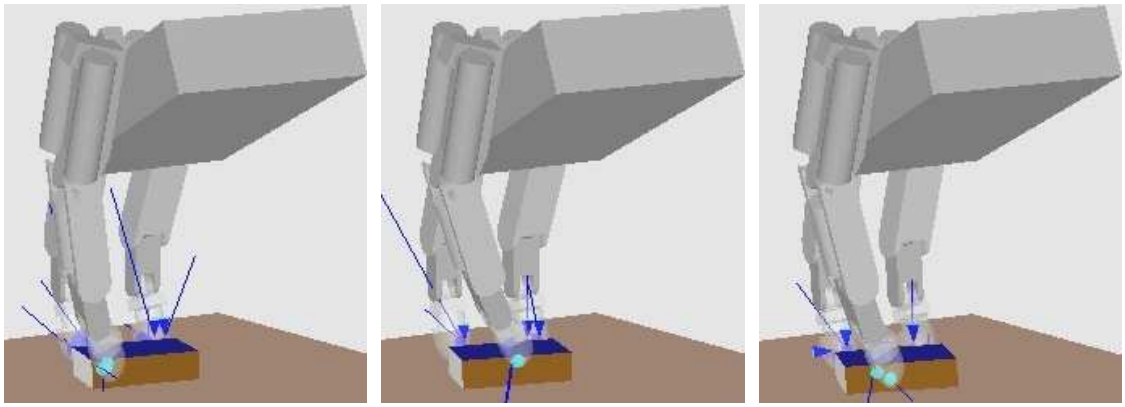

(a) *standard grasp:* thumb touches the object close to a corner. (b) *new pre-grasp:* thumb contact is near the upper edge. (c) *new target grasp:* thumb touches the object at a lower point.

Figure 7.11: TUM Hand grasping the matchbox in simulation.

- *standard grasp:* The thumb touches the object close to a corner of the side opposed to the other finger contacts (see Figure 7.11a). In three cases this side was missed in the evaluation with the real hand.

- *new pre-grasp:* As Figure 7.11b shows, the thumb touches the matchbox nearer its centre of mass when the *new pre-grasp* is utilised. However, the contact point is located at the upper edge, and a kind of fortune is responsible that the thumb did not glide over this edge in any real grasp trial.
- *new target grasp:* The new target grasp results in a contact point farther away from the centre of mass (see Figure 7.11b). Nevertheless, it is located lower and thus leads to a larger stability value. Grasping the matchbox with the real hand showed no uncertainty in nine out of ten trials. But again the thumb touches the object closer to the corner of the side. Hence, the grasp is less robust against disturbances. The variance in object position and orientation and the inaccuracies in the finger control of the TUM Hand are the reasons for that in one grasp trial the thumb misses that opposing side.

15. light bulb:

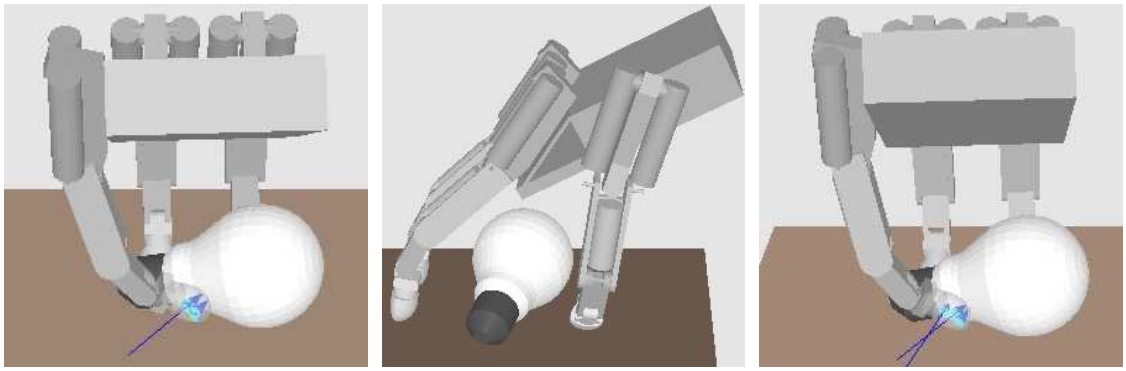

(a) *standard grasp:* thumb pushes the light bulb laterally.

(b) *new pre-grasp:* hardly any difference to the standard pre-grasp posture.

(c) *new target grasp:* optimal thumb closure trajectory.

Figure 7.12: TUM Hand grasping the light bulb in simulation.

- *standard grasp:* The thumb trajectory is suboptimal since it exerts forces that push the light bulb to the ulnar side of the hand (see Figure 7.12a). Therefore, the object is lost during the grasp simulation and the stability σ is low. When grasping with the real TUM Hand, it was also observed that the thumb pushed the light bulb laterally, and the object was lost in four of ten grasp trials.
- *new pre-grasp:* Due to the curved shape of the light bulb, the index and the middle fingers already touch the object simultaneously when using the standard grasp. Hence, there is no change in their pre-grasp posture after the first optimisation step (see Figure 7.12b). Additionally, the thumb is just slightly more flexed, but this difference leads to a small improvement in contact simultaneity and to a success rate of nine out of ten grasp trials.
- *new target grasp:* The new target grasp leads to an optimal thumb closure trajectory (see Figure 7.12c) that keeps the thumb from pushing the light bulb laterally. As a result, the object is not lost anymore in simulation and in reality.

16. chocolate bar:

- *new pre-grasp:* When using the *new pre-grasp*, the chocolate bar is less rotated, and it is touched by the thumb nearer its centre of mass (compare Figures 7.13a and 7.13b).

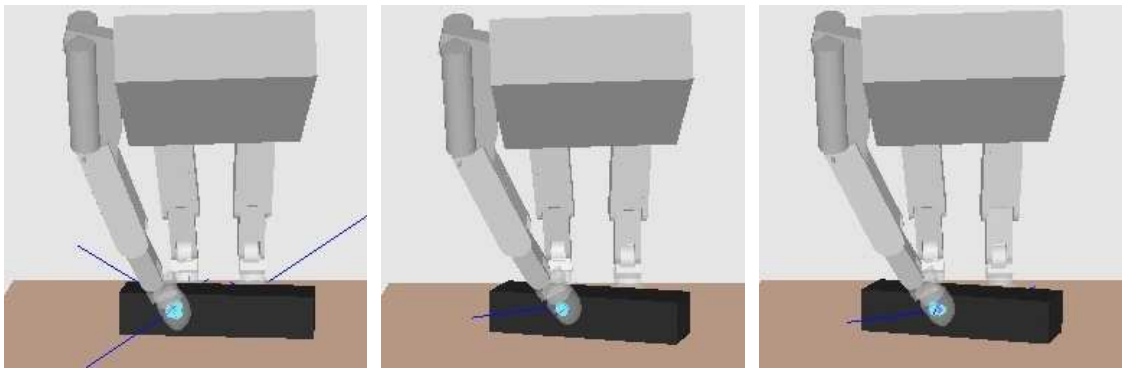

(a) *standard grasp:* object is rotated by the index finger (first contact).

(b) *new pre-grasp:* thumb touches object nearer its centre of mass.

(c) *new target grasp:* object is touched still nearer its centre of mass.

Figure 7.13: TUM Hand grasping the chocolate bar in simulation.

This leads to a stability σ larger than twice the value of the standard grasp and to full success in the grasp evaluation.

- *new target grasp:* The stability σ is enhanced again because the thumb touches the chocolate bar still nearer the centre of mass than before (see Figure 7.13c).

17. folding rule:

    - *new pre-grasp:* The optimisation of the pre-grasp suffers from the same effects as described in the case of the board marker. The object also rotated between the fingers and was lost in the lift-off phase when grasping with the real hand. But in contrast to the board marker, the folding rule was grasped in two trials successfully.
    - *new target grasp:* Optimising the thumb angles leads to an optimal thumb closure trajectory and to full success in the grasp experiment.

18. voltage tester:

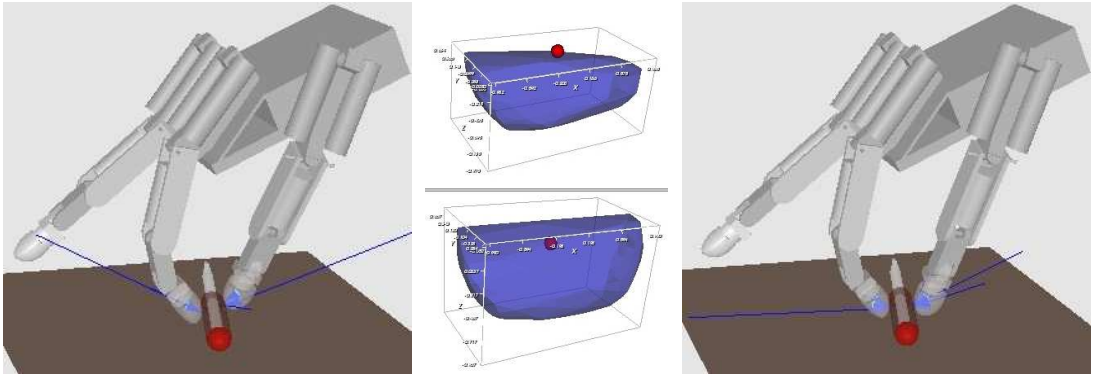

(a) *standard grasp:* grasp is not force closure at the moment of first touch of the fingers.

(b) force space: top: *standard grasp*; bottom: *new target grasp*.

(c) *new target grasp:* grasp is force closure, even with default friction coefficients.

Figure 7.14: TUM Hand grasping the voltage tester in simulation.

- *standard grasp:* The stability evaluation of the standard grasp results in a value of $\sigma < 1e^{-6}$. At the point in time when both fingers touch the voltage tester (see

Figure 7.14a), a force in direction of gravity cannot be resisted. This is proved by the force part of the wrench space shown at the top in Figure 7.14b. This uncertainty was also observed when grasping with the real hand while only three grasp trials were successful by chance.

- *new pre-grasp:* The new pre-grasp posture does not lead to essential differences when the fingers touch the voltage tester. Similar as before, the thumb nearly reaches its target grasp posture when it touches the object. Hence, the stability values and the grasp success rates are comparably low.
- *new target grasp:* For optimising the thumb target grasp posture, the friction coefficients between the materials of the fingertips and the voltage tester were doubled ($\mu = 2.0$; $\nu = 0.6$; see Section 6.1.2). This offered the possibility to achieve an optimal thumb posture and a stable grasp in simulation. Even when using the default friction coefficients (see Figures 7.14b (bottom) and 7.14c), a stability of $\sigma = 0.065514$ is achieved, whereas the value listed in Table 7.1 corresponds to high friction. When evaluating the optimised grasp in reality, only in one trial the object was lost.

19. eraser:

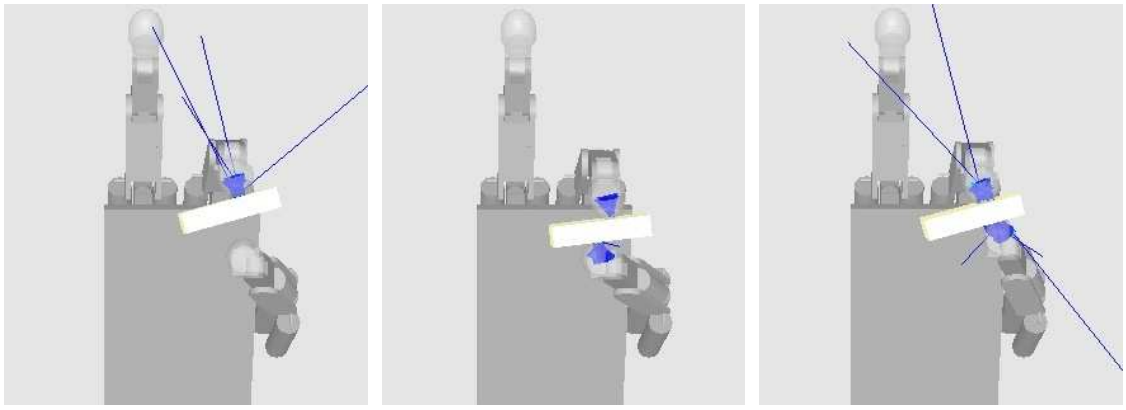

(a) *standard grasp:* index finger drags and rotates the object. (b) *standard grasp:* suboptimal opposition of finger contacts. (c) *new target grasp:* both fingers touch the object simultaneously

Figure 7.15: TUM Hand grasping the eraser in simulation.

- *standard grasp:* The index finger touches the object first and drags it towards the thumb (see Figure 7.15a). Because the contact is not near the object's centre of mass, the eraser it rotated. The grip posture does not provide a suitable finger opposition to the thumb (see Figure 7.15b). This leads to a low stability σ and is the reason why in most grasp trials executed with the real hand the eraser was lost.
- *new pre-grasp:* Since, in the case of the voltage tester, the *new pre-grasp* only leads to small differences in the grip posture, the stability values are similar and the grasp success with the real hand was larger only by chance.
- *new target grasp:* This optimisation step leads to a thumb posture with a more suitable opposition to the index finger than before (see Figure 7.15c). Additionally, now both fingers touch the eraser at the same time. The evaluated stability σ is much higher, and in reality no grasp trial was incomplete.

20. bunch of keys:

    - *new pre-grasp:* The bunch of keys is modelled as a simple box in simulation (see Section 6.1.2). This is only a coarse approximation of this form-variable object, and the stability $\sigma > 0$ does not correspond to reality. Grasping with the real TUM Hand was not successful in any trial both before and after the first optimisation step.
    - *new target grasp:* The thumb angles of the target grasp only change slightly when applying the second optimisation step. Again, the experiment showed that the three-fingered TUM Hand is hardly able to grasp this form-variable object.

21. pencil:

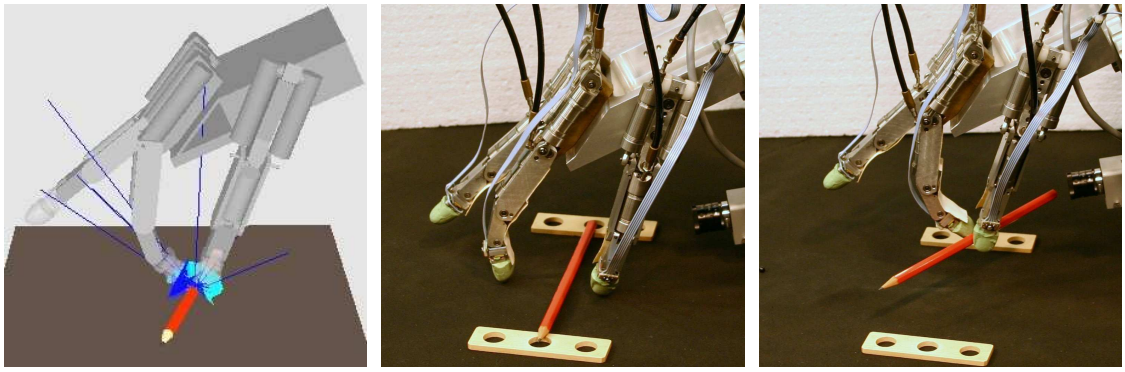

(a) *new target grasp:* high friction coefficients lead to wide friction cones. (b) *new target grasp:* pads elevate the pencil by 4 mm. (c) *new target grasp:* successful grasp of elevated pencil.

Figure 7.16: TUM Hand grasping the pencil in simulation and reality.

- *standard grasp:* The standard two finger precision grasp is not appropriate for grasping the pencil. In simulation the thumb does not even touch this thin object, thus the stability $\sigma = 0$.
- *new pre-grasp:* The *new pre-grasp* does not lead to an improvement since the grasp posture of the thumb is responsible for that the pencil cannot be grasped.
- *new target grasp:* The stability of the optimised grasp is low ($\sigma < 1e^{-6}$) when using realistic friction coefficients in simulation. For optimising, four times larger coefficients were used ($\mu = 4.0$; $\nu = 1.2$) leading to wide friction cones (see Figure 7.16a) and to great stability σ as listed in Table 7.1. The new target grasp realises the most appropriate two finger precision grasp for grasping the pencil with the real TUM Hand, but the pencil was not lifted in any of the ten trials. The reason is a remaining gap between the tips of the thumb and the index finger in target grasp posture. Nevertheless, the fact that the optimised target grasp is most suitable for grasping this object was shown after elevating the pencil by 4mm using two pads (see Figure 7.16b). This led to ten out of ten successful grasp trials in an additional experiment (see Figure 7.16c).

7.1.2 Simulated and Real Grasps of the Shadow Hand

no.	name	grasp type	standard grasp		new pre-grasp		new target grasp	
			stability σ	success	stability σ	success	stability σ	success
1	adhesive tape	t_3	0.157609	10	0.185234	10	0.337551	10
2	toy propeller	t_5	0.172909	10	0.234636	10	0.292682	10
3	toy cube	t_4	0.248465	10	0.248465	(10)	0.380063	10
4	can	t_3	0.255011	10	0.295776	10	0.382037	10
5	tissue pack	t_3	0.183172	10	0.137026	10	0.240717	10
6	tennis ball	t_3	0.114489	7	0.312271	8	0.362121	10
7	paper ball	t_3	0.219402	10	0.222263	10	0.300357	10
8	sharpener	t_1	0.257970	10	0.303304	10	0.424608	10
9	remote control	t_3	0.284305	10	0.127224	10	0.417855	10
10	cup	t_3	0.138124	10	0.431006	–	0.455679	10
11	board marker	t_2	0.030594	10	0.034759	10	0.056605	10
12	tea light	t_1	0.091017	8	0.241182	9	0.268954	10
13	golf ball	t_3	0.025039	6	0.101005	8	0.224461	9
14	matchbox	t_1	0.238027	6	0.252589	9	0.335330	10
15	light bulb	t_3	0.017925	8	0.021799	10	0.342233	10
16	chocolate bar	t_1	0.246085	10	0.248558	10	0.274517	10
17	folding rule	t_2	0.016603	10	0.015223	10	0.110728	10
18	voltage tester	t_2	0.017331	8	0.017107	7	0.080589	9
19	eraser	t_2	0.054965	9	0.051065	10	0.241145	1/10
20	bunch of keys	t_1	0.232630	1	0.211252	0	0.260685	2
21	pencil	t_2	0	0	0	0	0.122870	8

Table 7.2: Grasp stability σ of simulated grasps and the number of successful grasp trials (0 to 10) out of 10 grasp attempts with the real Shadow Hand before optimisation (standard grasp), after the first optimisation step (new pre-grasp), and after the second optimisation step (new target grasp). Red values indicate *new pre-grasps* that are used in the second optimisation step.

Evaluation of the experiments on grasping the benchmark objects with the *standard grasp*, after the first optimisation step (*new pre-grasp*), and after the second optimisation step (*new target grasp*):

1. adhesive tape:

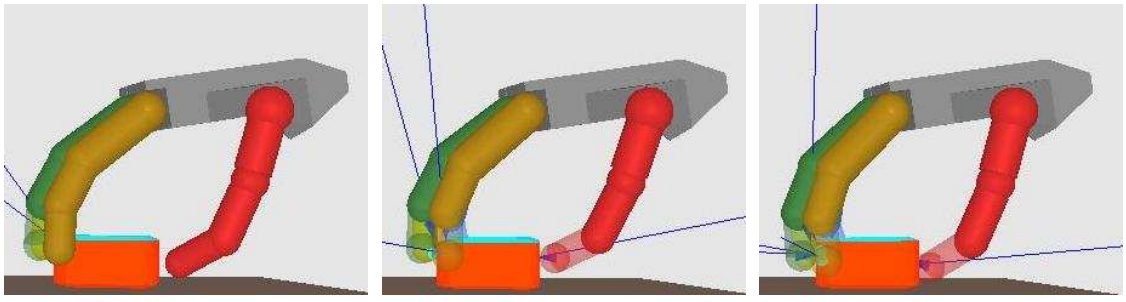

(a) *standard grasp:* opposing fingers "drag" the object towards the thumb. (b) *new pre-grasp:* "both" opposing forces are applied simultaneously. (c) *new target grasp:* thumb touches object at a lower point.

Figure 7.17: Shadow Hand grasping the adhesive tape in simulation.

- *standard grasp:* In some of the ten grasp trials, the adhesive tape rotated between the fingers after it was grasped. This was due to the fact that no force sensors were used

(see Section 3.1.3) and the finger controller continually tried to reach the target grasp angles for all joints. Nevertheless, the adhesive tape was grasped successfully in each trial.

- *new pre-grasp:* The object rotation observed previously was reduced and in some grasp trials totally avoided by using the optimised pre-grasp which converts the *finger drag* strategy (see Figure 7.17a) of the standard grasp into a *both* strategy (see Figure 7.17b). The finger contact simultaneity leads to an equalisation of torques exerted on the object and thus to less rotation.
- *new target grasp:* The optimised target grasp posture results in a contact at a lower and more centred point upon the surface of the adhesive tape (see Figure 7.17c). Because the thumb provides the major part of the opposing force to each force exerted by the fingers, an optimised opposition posture of the thumb can lead to a much larger stability value, as it is the case when grasping this object in simulation. The experiment with the real Shadow Hand revealed no essential differences when evaluating the grasps after each optimisation step.

2. toy propeller:

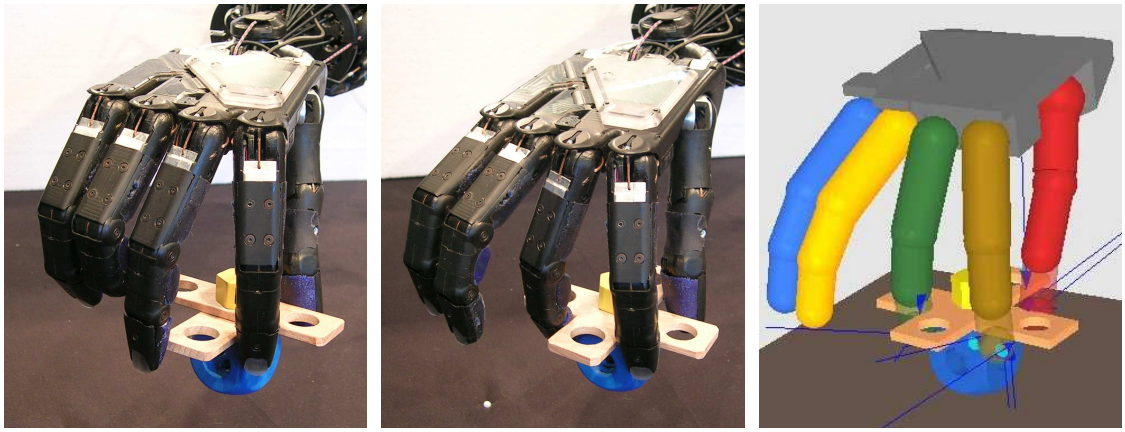

(a) *standard grasp:* middle finger touches propeller at its blade.

(b) *new pre-grasp:* fingers reach between the blades of the toy propeller.

(c) *new target grasp:* simulated three finger special grasp.

Figure 7.18: Shadow Hand grasping the toy propeller in reality and in simulation.

- *standard grasp:* Grasping the toy propeller revealed a major difference between the TUM and the Shadow Hand concerning the configuration of the fingers at the palm. Because the distances between the metacarpophalangeal joints of the fingers are shorter in the case of the Shadow Hand, it happened that one finger got stuck at a blade of the propeller (see Figure 7.18a) rather than touching the object between the blades (see Figure 7.18b). But even when touching the tip of the blade, the grasps were successful.
- *new pre-grasp:* Since the standard pre-grasp and the optimised pre-grasp just differ slightly, no essential differences were observed comparing both grasp experiments. Again, all grasp trials were successful.
- *new target grasp:* The optimised thumb joint angles are similar to those of the standard grasp. Nevertheless, small differences lead to a larger stability value in simulation (see Figure 7.18c).

3. toy cube:

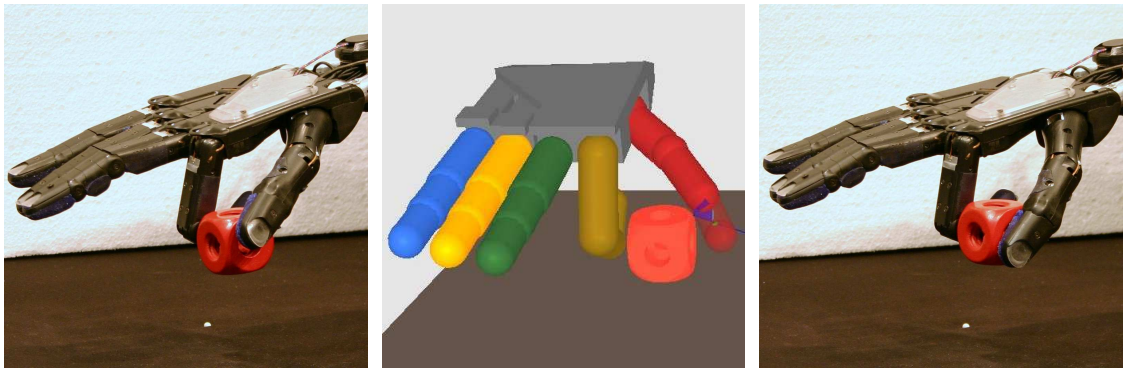

(a) *standard grasp:* thumb touches the toy cube at its upper edge.

(b) *standard grasp:* index finger reached its target grasp posture before the thumb touches the object.

(c) *new target grasp:* thumb touches the toy cube at a lower point.

Figure 7.19: Shadow Hand grasping the toy cube in reality and in simulation.

- *standard grasp:* Although the toy cube was successfully grasped in each trial, it sometimes was touched by the thumb at the upper edge (see Figure 7.19a). Hence, more uncertainties in object location could easily lead to an incomplete grasp.
- *new pre-grasp:* The first optimisation step does not result in different pre-grasp joint angle values. The reason is that the index finger reached its target grasp posture already when the thumb touches the toy cube as shown in Figure 7.19b (compare Section 6.3.1). Since there is no difference compared to the standard grasp, the stability values are equal and a re-evaluation of the same grasp was not performed. The brackets around the number of successful grasp trials in Table 7.2 denote that the success rate of the standard grasp was assumed.
- *new target grasp:* The optimised thumb joint angle values lead to great stability σ caused by a thumb contact at a lower point of the toy cube. Although the success rate does not differ to that of the standard grasp, the optimised grasp applied with the real hand shows no uncertainty anymore (see Figure 7.19c).

4. can:

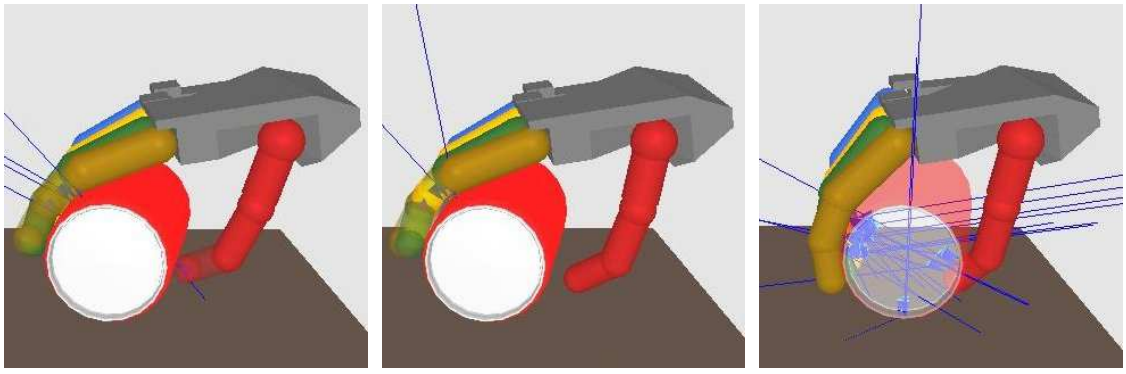

(a) *new pre-grasp: both* contact strategy is achieved.

(b) *new target grasp:* first contact in the *finger drag* strategy.

(c) *new target grasp:* enclosing grasp at first contact of all fingers.

Figure 7.20: Shadow Hand grasping the can in simulation.

- *standard grasp:* The can was grasped in each of the ten trials successfully, but it was observed that there is only small variability in the object's position and orientation relative to the hand being in grasp position. The values of these parameters have to match exactly those assumed by the robot system because of the large size of the object. Otherwise the fingers cannot enclose the object and it could – initiated by the fingers in pre-grasp posture – roll away before the grasp is applied.
- *new pre-grasp:* Due to the fact that the shape of the pre-grasp posture now is more adopted to the shape of the can, the object is less rotated before it is lifted. The thumb is more flexed for achieving finger contact simultaneity (see Figure 7.20b) resulting in a smaller hand opening. This leads to less variability in the positioning of the hand, but the can again was grasped successful in each of the ten trials.
- *new target grasp:* With the second optimisation step, the *both* contact strategy achieved with the optimised pre-grasp is changed into a *finger drag* strategy again. The thumb closure trajectory ends at a location that is farther away from the object and the other fingertips than before (see Figure 7.20b). Hence, the object is dragged towards the thumb and the results is a more object enclosing grip posture and a larger stability σ (see Figure 7.20c).

5. tissue pack:

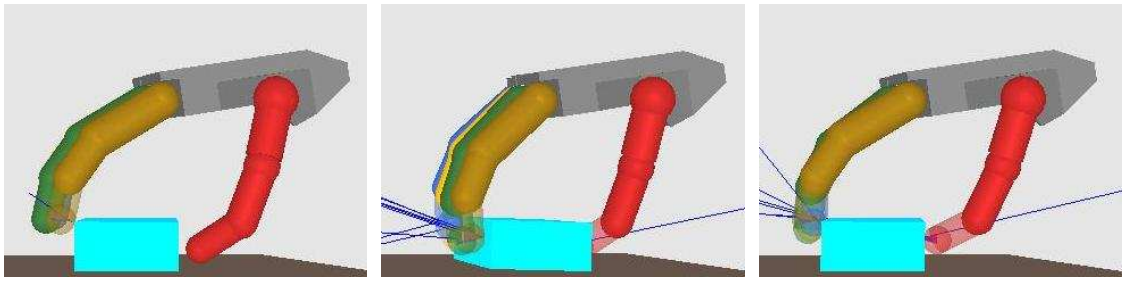

(a) *standard grasp:* first finger contact. (b) *standard grasp:* first moment of all finger contacts. (c) *new pre-grasp:* first moment of all finger contacts.

Figure 7.21: Shadow Hand grasping the tissue pack in simulation.

- *standard grasp:* As the simulation of the standard grasp shows (compare Figures 7.21a and 7.21b), the object is rotated during grasp closure. The same effect was observed when grasping with the real hand. But the form-variable and compliant tissue pack is formed to the fingers when they exert pressure on it, and this enhances the chance of grasp success. Hence, the object was grasped successfully in all trials.
- *new pre-grasp:* The one-shot learning step for optimising the pre-grasp posture leads to a less stable grasp although it realises a perfect finger contact simultaneity. The *both* contact strategy is achieved particularly by a more flexed thumb in the optimised pre-grasp posture. Thus, the fingers are more extended at the moment when the thumb touches the object (compare Figures 7.21b and 7.21c). As a consequence, the fingers touch the tissue pack nearer its upper edge leading to a lower stability σ. But because of the advantageous properties of the object for grasping it, each grasp trial was successful.
- *new target grasp:* Grasping the tissue pack with the optimised thumb joint angle values revealed no essential differences than when using the standard grasp. Again, each of the ten grasp trials was successful.

6. tennis ball:

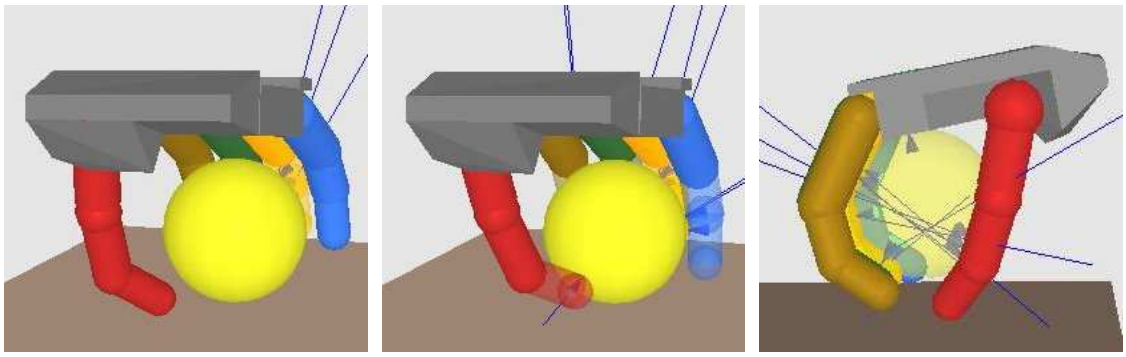

(a) *standard grasp:* ring finger touches object first.

(b) *new pre-grasp:* enclosing pre-grasp leads to simultaneous contacts.

(c) *new target grasp:* ”optimal” power grasp involving the palm.

Figure 7.22: Shadow Hand grasping the tennis ball in simulation.

- *standard grasp:* In three grasp trials the tennis ball was pushed out of the hand after being grasped successfully. It was observed that the index, the middle, and the ring fingers exerted too much force that the thumb could not resist. The reasons are a suboptimal thumb closure trajectory and the fact that no touch sensors were utilised for avoiding further motion of the fingers after touching the object.
- *new pre-grasp:* The more object enclosing pre-grasp (compare Figures 7.22a and 7.22b) inhibits the rolling of the object out of the hand. But still in two grasp trials the object was lost because of the same reasons responsible for the incomplete grasp trials in the case of the standard grasp.
- *new target grasp:* After optimising the target grasp, all grasp trials were successful. In the optimal opposition posture, the thumb resists all forces exerted by the fingers on the tennis ball, and an ”optimal” power grasp is accomplished by pressing the object against the palm (see Figure 7.22c).

7. paper ball:

- *standard grasp:* The advantage of the Shadow Hand having more fingers compared to the TUM Hand especially comes to grip when objects like the form-variable paper ball have to be grasped. Whereas, in the case of the TUM Hand, the deformation of the paper ball led to incomplete grasp trials in some cases, when grasping with the Shadow Hand, the object’s shape is adapted to the five grasping fingers leading to successful grasp trials.
- *new pre-grasp:* As also indicated by the stability σ, the *new pre-grasp* is not advantageous for grasping the paper ball compared to the standard pre-grasp. Although now all fingers touch the paper ball simultaneously in simulation, no essential differences were observed when grasping with the real hand.
- *new target grasp:* The differences of the thumb angle values are less than three degrees in each joint compared to those of the standard grasp. When grasping the form-variable paper ball with the real hand, this order of magnitude does not result in remarkable changes. Again, the object is adapted to the five fingers of the Shadow Hand leading to successful grasp trials.

8. sharpener:

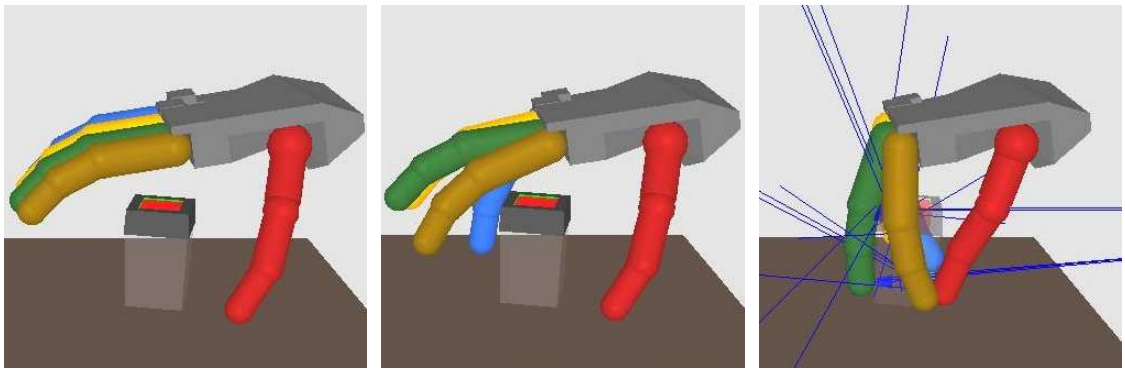

(a) *standard grasp:* standard pre-grasp posture.

(b) *new pre-grasp:* optimised pre-grasp posture.

(c) *new target grasp:* high stability through enclosing fingers.

Figure 7.23: Shadow Hand grasping the sharpener in simulation.

- *standard grasp:* When grasping with the five-fingered Shadow Hand, it is not that severe if one finger misses the object or if one finger touches the object at an edge, compared to grasping with the three-fingered TUM Hand. Although the sharpener was pushed and rotated in different directions during the execution of the grasp trials, its motion always was completed by a certain grip posture.
- *new pre-grasp:* The *new pre-grasp* shows hardly any difference in the thumb angles, but particularly the index finger and the little finger better enclose the object (compare Figures 7.23a and 7.23b). This leads to a grasp with larger stability σ and to less movement of the sharpener in the grasp trials executed with the real hand.
- *new target grasp:* The magnitude of the stability σ after optimising the target grasp is comparable to that of object enclosing power grasps. In this all finger precision grasp, not only the sharpener is touched with the fingertips, but also contacts with other phalanges occur (see Figure 7.23c). The new target grasp posture of the thumb causes the object to move even less during the grasp process than it was observed after the first optimisation step.

9. remote control:

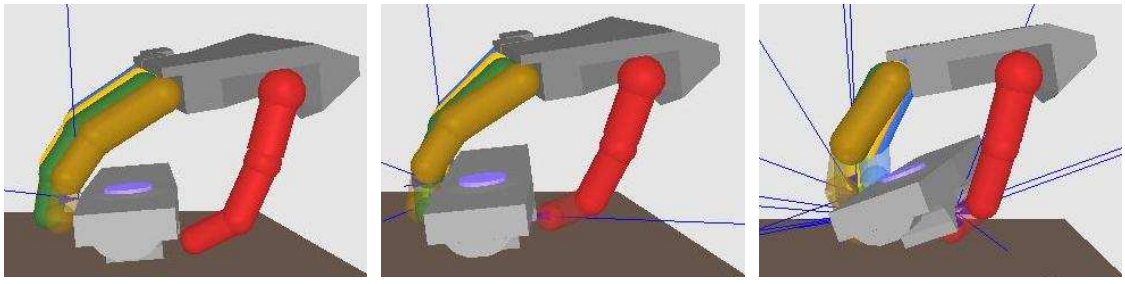

(a) *standard grasp:* index finger sets remote control on the edge.

(b) *new pre-grasp:* thumb touches remote control at its side.

(c) *new target grasp:* thumb matches the bottom shape of the remote control.

Figure 7.24: Shadow Hand grasping the remote control in simulation.

- *standard grasp:* In simulation, the remote control is set on the edge (see Figure 7.24a) before the thumb reaches under its bottom side, leading to an enveloping grasp. The stability σ of this power grasp is larger than those of any other standard grasp applied with the Shadow Hand. This enveloping effect was not observed when grasping with the real hand, where the thumb touched the object rather at the side than at its bottom. This discrepancy is due to simplifying the model of the remote control in the simulator by assuming sharp edges.
- *new pre-grasp:* When using the optimised pre-grasp in simulation, now the thumb touches the remote control at its side during grasp closure (see Figure 7.24b). In contrast to the standard grasp, the thumb does not reach under the remote control anymore. The object is not enveloped and thus the stability σ is lower. When grasping with the real hand, no essential differences were observed, and each grasp trial was successful again.
- *new target grasp:* Because the stability value was not enhanced by the first optimisation step, the standard pre-grasp was utilised for optimising the joint angle values of the thumb. When simulating the grasp after this second optimisation step, the enveloping effect occurs again. The thumb matches the bottom shape of the remote control exactly (see Figure 7.24c). This leads to a large stability σ.

10. cup:

(a) *standard grasp:* tip of the thumb touches the cup. (b) *new pre-grasp:* thumb is flexed to far in pre-grasp posture. (c) *new target grasp:* thumb matches shape of the cup.

Figure 7.25: Shadow Hand grasping the cup in simulation.

- *standard grasp:* Similar to the can, there is only small variability in the grasp position relative to the large-sized cup. An inaccurate positioning of the real robot hand above the object leads to contacts between the fingers in pre-grasp posture and the cup during the placing phase. Since this was not observed when evaluating the standard grasp (see Figure 7.25a), all grasp trials were successful.
- *new pre-grasp:* When optimising the pre-grasp for grasping the cup, the same problem as in the case of the TUM Hand occurred. Again, the distance between the tip of the thumb and the cup is too small in the optimised pre-grasp posture (see Figure 7.25b). Although there is no contact between the thumb and the cup when the hand is in grasp position, the optimised pre-grasp posture of the thumb would lead to a contact with the seam of the cup during the placing phase. Thus, an evaluation of the *new pre-grasp* was not performed.

- *new target grasp:* Because the optimised pre-grasp is not applicable, the standard pre-grasp was utilised for optimising the target grasp. Nevertheless, the resulting grasp leads to a large stability σ. When applying the new target grasp, the thumb matches the shape of the cup which leads to more contact points than before (see Figure 7.25c). The grasp trials executed with the real hand showed no uncertainty.

11. board marker:

    - *standard grasp:* Since only two fingers are involved when grasping the board marker (two finger precision grasp), these two fingers have to be in exact alignment for providing the opposing forces required. The rate of ten successful grasps trials shows that the standard two finger precision grasp is appropriate for grasping the board marker.
    - *new pre-grasp:* The changes in the joint angle values are small between the standard grasp and the *new pre-grasp*. Hence, the corresponding stability values just differ slightly and the board marker was grasped successfully in each trial.
    - *new target grasp:* The second optimisation step leads to a change below two degrees in each of the thumb joint angle values. This leads to no essential differences when evaluating the optimised grasp with the real hand.

12. tea light:

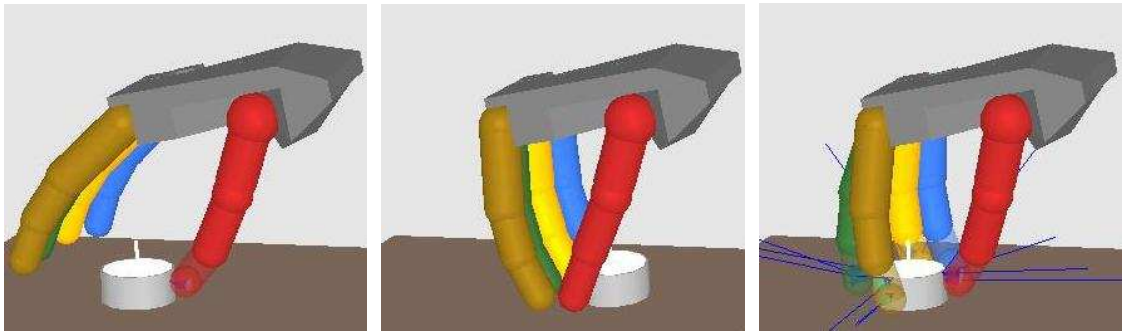

(a) *standard grasp:* thumb touches object first (*thumb push* strategy). (b) *standard grasp:* object is lost after being grasped. (c) *new target grasp:* grasp after optimising in two steps.

Figure 7.26: Shadow Hand grasping the tea light in simulation.

- *standard grasp:* Because the tea light is a small object with low frictional coefficients and is roundish at its contact surface, it was lost after being grasped successfully in two out of ten grasp trials. This effect is also shown when simulating the standard grasp (see Figure 7.26b) and results in a small stability σ.
- *new pre-grasp:* After optimising the pre-grasp, a *both* contact strategy is realised. In contrast to the prior *thumb push* strategy (see Figure 7.26a), large motions of the object before being grasped are inhibited. But still the object was pushed out of the hand in one grasp trial. In simulation, the tea light is not lost anymore leading to a large improvement of the stability σ.
- *new target grasp:* The optimised thumb closure trajectory endows the thumb with the ability to resist all forces exerted by the opposing fingers when touching the object (see Figure 7.26c). The results are a larger stability σ and full success in the grasp experiment with the real hand.

13. golf ball:

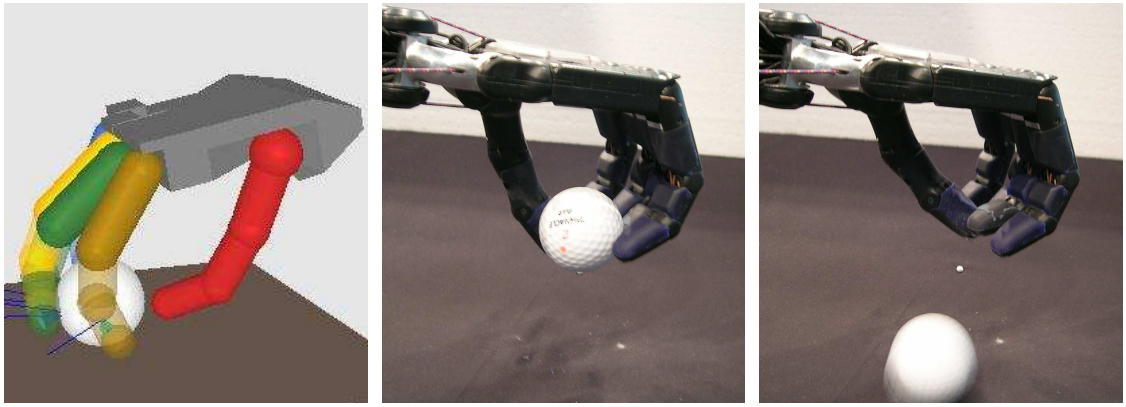

(a) *new pre-grasp:* object enclosure inhibits lateral motion. (b) *new target grasp:* object is grasped at first... (c) *new target grasp:* ...and is lost afterwards.

Figure 7.27: Shadow Hand grasping the golf ball in simulation and in reality.

- *standard grasp:* The golf ball resembles the tea light in its properties relevant for grasping. Hence similar problems occur like a low stability σ and the fact that in some grasp trials the object was pushed out of the hand. But in contrast to the tea light, there is less variability in the position of the golf ball relative to the grasping hand. An inexact grasp position leads to an earlier contact with a single finger causing the golf ball to roll away before all grasping fingers touch it.
- *new pre-grasp:* The optimisation of the pre-grasp leads to an enclosure of the golf ball when the fingers (except for the thumb) touch the object simultaneously. This inhibits the lateral movement of the golf ball (see Figure 7.27a), but it is rolled towards the thumb before being grasped (*finger drag* strategy). Nevertheless, this rolling is better guided by the enclosing fingers leading to more successful grasp trials.
- *new target grasp:* By utilising the results of the second optimisation step, the grasp success rate was improved again. Unfortunately, one grasp trial was incomplete although the golf ball was successfully grasped at first (see Figure 7.27b). But because the index and the middle finger exerted too much force, the object was pushed out of the grasping hand afterwards (see Figure 7.27c). This problem can be avoided by using fingertip sensors to stop the fingers from further motion.

14. matchbox:

- *standard grasp:* The evaluation of grasping the matchbox evidenced the fact that for flat objects, it is important that the fingers close directly above the desktop surface (while avoiding to stick into it). Although the standard all finger grasp meets this requirement, the matchbox was not grasped in four of the ten grasp trials. In the incomplete trials, the object was rotated and the thumb slipped over its upper edge (see Figure 7.28a).
- *new pre-grasp:* Before optimising the pre-grasp, the thumb touched the matchbox first (see Figure 7.28b). When applying the optimised pre-grasp, all grasping fingers touch the object simultaneously (see Figure 7.28c) inhibiting a rotation of the object. Hence, a slip of the thumb over the flat object in less likely, and all but one grasp trials were successful.

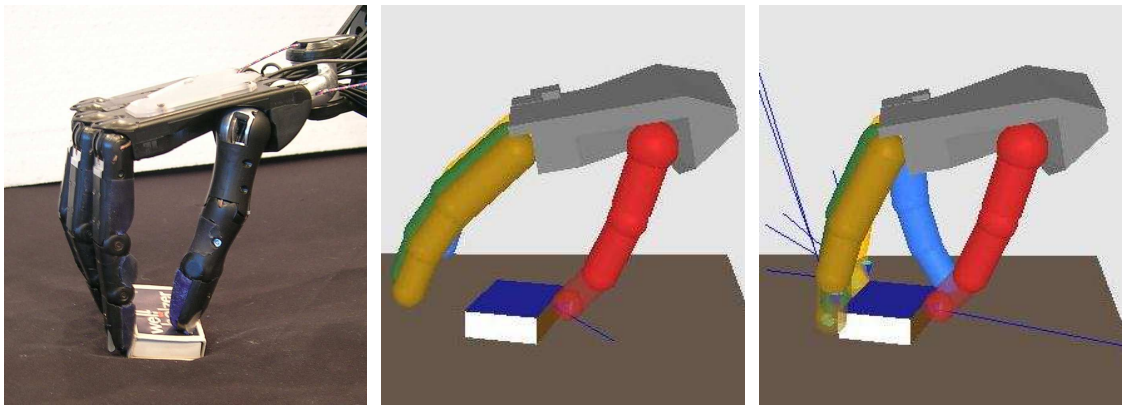

(a) *standard grasp:* thumb glides over the object. (b) *standard grasp:* first touch (*thumb push* strategy). (c) *new pre-grasp:* simultaneous contacts (*both* strategy).

Figure 7.28: Shadow Hand grasping the matchbox in reality and in simulation.

- *new target grasp:* Optimising the thumb closure trajectory led to a large stability σ in simulation, and no grasp trial with the real hand was incomplete anymore.

15. light bulb:

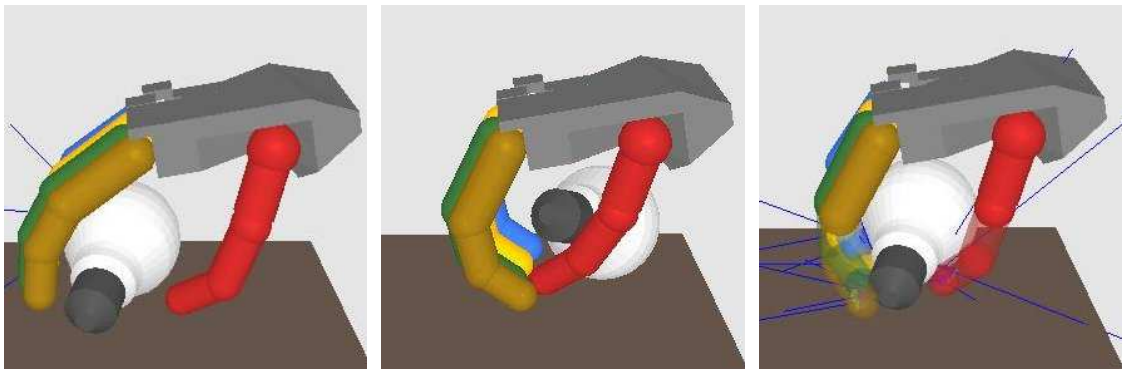

(a) *standard grasp:* object is dragged by first fingers touching it. (b) *standard grasp:* object is lifted over the thumb. (c) *new target grasp:* grasp optimised in two steps is successful.

Figure 7.29: Shadow Hand grasping the light bulb in simulation.

- *standard grasp:* The ring finger and the little finger dragged the light bulb and lifted it over the thumb in two out of ten grasp trials. Additionally, the closing thumb caused a lateral motion of the object and supported its slip out of the hand. The same effects are observed in simulation (see Figures 7.29a and 7.29b).
- *new pre-grasp:* After the first optimisation step, all fingers, except for the thumb, touch the light bulb simultaneously preventing the object from being lifted over the thumb. The second improvement of the optimised pre-grasp is that the thumb nearly reached its target grasp posture when touching the object. Hence, the thumb does not cause a lateral motion of the object anymore, and all grasp trials were successful when grasping with the real hand. In simulation, the light bulb is still lost after being grasped successfully. Thus, the stability σ is enhanced only slightly.
- *new target grasp:* The second optimisation step leads to a much larger stability σ because the object is not lost anymore in simulation (see Figure 7.29c). The optimised

thumb trajectory ends in a more distal point leading to a contact with the object occurring earlier in the grasp closure phase. Since the first optimisation step led to full grasp success already, in the experiment of grasping with the real hand after the second step no essential differences were observed.

16. chocolate bar:

    - *standard grasp:* When applying the standard all finger precision grasp, the first finger touching the chocolate bar is the index finger. This finger exerts a torque on this oblong object by which it is slightly rotated before being touched by all grasping fingers. Nevertheless, the object was grasped successfully in each grasp trial.
    - *new pre-grasp:* With the first optimisation step, contact simultaneity of all fingers, except for the thumb, is achieved. This prevents the object from being rotated and again each grasp trial was successful.
    - *new target grasp:* The optimised thumb trajectory realises a contact with the chocolate bar nearer the object's centre of mass and results in a larger stability σ.

17. folding rule:

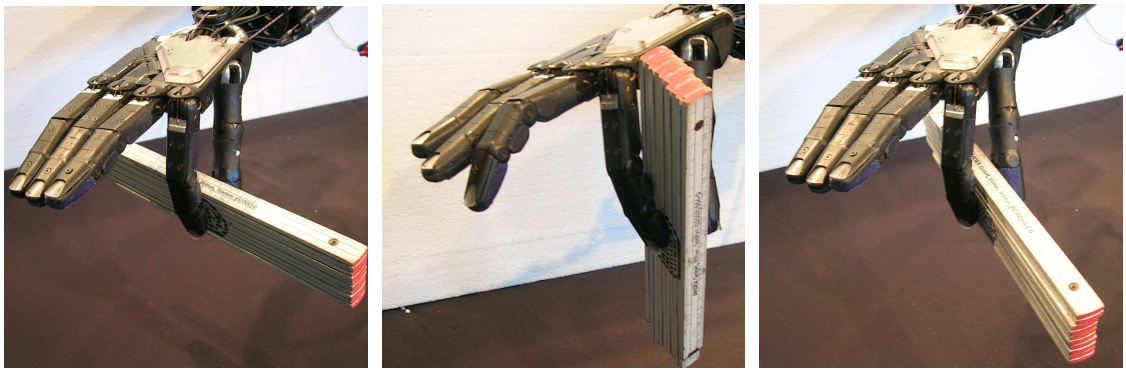

(a) *standard grasp:* tweezers-like grasp (in balance). (b) *standard grasp:* object rotates between fingers (c) *new target grasp:* pinch-like grasp.

Figure 7.30: Shadow Hand grasping the folding rule in reality.

- *standard grasp:* The folding rule has to be touched near its centre of mass when being grasped with a two finger precision grasp (see Figure 7.30a). Otherwise the object rotates between the thumb and the index finger in the lift-off phase. The reason is that the mass distribution of this relatively heavy and longish object causes an imbalance that the forces exerted by the two fingers could not resist. This happened in some grasp trials (see Figure 7.30b), but the object was lost in no case.
- *new pre-grasp:* The differences in the joint angle values in relation to the standard pre-grasp are negligible. Thus, the stability values are similar and all grasp trials were successful again.
- *new target grasp:* Because the optimised thumb trajectory is directed more towards the radial side of the index finger, the folding rule is pressed rather against the side of the index fingertip than against its volar pad. Hence, the former tweezers-like two finger precision grasp now becomes a more pinch-like grasp (compare Figures 7.30a and 7.30c) with a larger stability σ.

18. voltage tester:

    - *standard grasp:* The grasp experiment with the real hand reveals that the standard two finger precision grasp is suitable for grasping the voltage tester. Compared to the board marker, the handle of the voltage tester is thinner and its shape is more complex. Hence, it is more difficult to grasp this object, and two out of ten grasp trials were incomplete.
    - *new pre-grasp:* The first optimisation step only leads to slight changes in the pre-grasp. The stability σ is similar to that of the standard grasp and only by chance one more grasp trial was incomplete.
    - *new target grasp:* The change in the optimised thumb joint angle values are only small, too. Nevertheless, it was observed that they lead to less object movement in the stabilisation phase than before. This is the reason why the grasp success was enhanced to nine out of ten grasp trials. In the incomplete trial, the two fingers touched the thinnest part of the object's handle.

19. eraser:

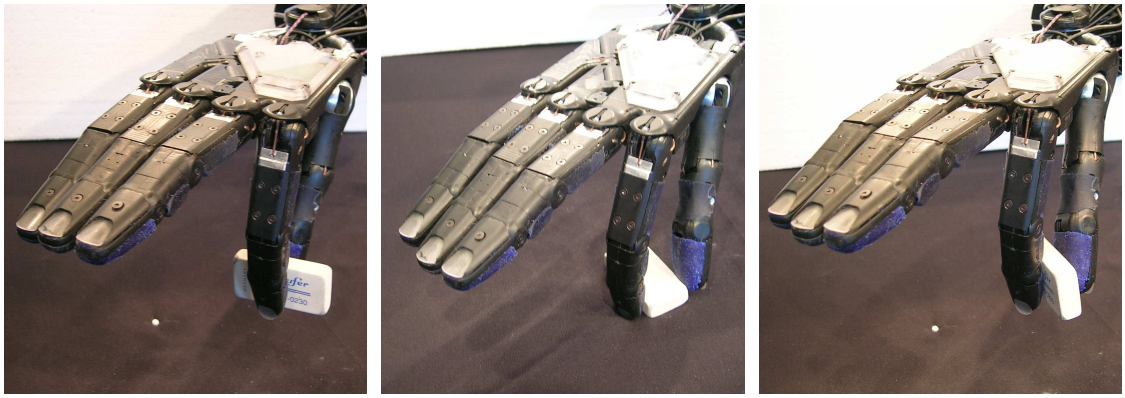

(a) *standard grasp:* tweezers-like grasp.

(b) *new target grasp (1):* object is knocked over by the thumb moving too fast.

(c) *new target grasp (2):* pinch-like grasp with correct finger control.

Figure 7.31: Shadow Hand grasping the eraser in reality.

    - *standard grasp:* The eraser is even thinner than the handle of the voltage tester, but the standard two finger precision grasp is also appropriate for successfully grasping this object (see Figure 7.31a). Only in one grasp trial the eraser was lost after being grasped by the two fingers. This was due to the fact that the finger joint angle controller tries to correct small deviations of the actual angles from the target angles. Caused by this active control, it happens that the sum of the joint angle deviations of both fingers result in a too large distance between the fingertips. This problem can be avoided by using fingertip sensors to stop the controller if contact with the object is detected.
    - *new pre-grasp:* Although the stability value is slightly lower, the grasp success rate is higher than that of the standard grasp. Like in the cases of the folding rule and the voltage tester, the optimised two finger precision pre-grasp only shows small changes. Hence, the differences in both the stability value and in the grasp success occurred by chance.
    - *new target grasp:* Similar to the second optimisation step in the case of grasping the folding rule, the two finger precision grasp turns into a pinch grasp by utilising the optimised target grasp. The thumb closure trajectory is directed more towards the radial

side of the index finger leading to an almost four times larger stability σ. In contrast, only a success rate of one out of ten grasp trials was achieved in the first evaluation of the optimised target grasp. We observed that the eraser was knocked over by the thumb before the index finger reached the object (see Figure 7.31b). The difference in the finger control between the simulated hand and the real Shadow Hand was due to frictional problems in the joints of the index finger, which occurred after an operation time of some weeks. Therefore, the initial mapping of the finger control into the simulation environment was not accurate any longer. After adapting the simulated finger joint velocities to that of the slowed down counterparts of the real hand, another attempt of optimising the thumb trajectory led to full success in the subsequent evaluation of the resulting target grasp (see Figure 7.31c). The stability $\sigma = 0.198909$ is almost three times larger than that of the standard grasp. Although the two grasping fingers apply a pinch-like grasp again (because the object is pressed by the thumb to the side of the index fingertip), in none of the grasp trials the eraser was knocked over anymore.

20. bunch of keys:

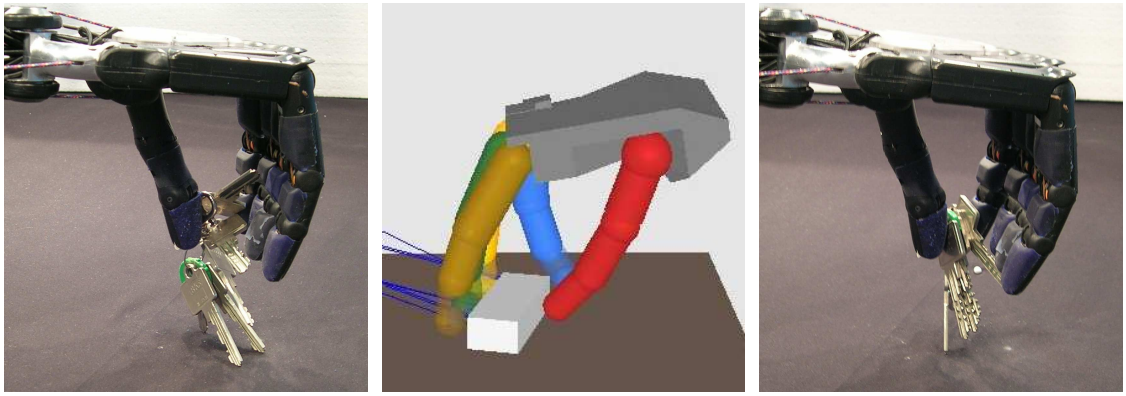

(a) *standard grasp:* opposing forces on one key.

(b) *new pre-grasp:* simultaneous contacts avoid rotation.

(c) *new target grasp:* several keys in vertical orientation.

Figure 7.32: Shadow Hand grasping the bunch of keys in reality and in simulation.

- *standard grasp:* The bunch of keys is the most difficult object to grasp. It consist mainly of five keys, and when it lies on the desktop, the outer keys slope down to the desktop surface (for illustration see Figure 4.1). Fingers, trying to grasp the bunch of keys, most often touch the flat keys at an edge and slip over their plain surface. Even if a finger adheres to a key, the keys move against each other enhancing the difficulty for exerting appropriate forces on the bunch. Only one out of ten grasp trials was successful, in which opposing forces have been exerted on the same key by chance (see Figure 7.32a). During the other grasp trials, the keys were pushed against each other. When the form-variable bunch of keys reaches a block-like alignment with its keys lying upon each other in a horizontal or a vertical (like shown in Figure 7.32c) orientation, there is a chance that the object is grasped successfully.
- *new pre-grasp:* The stability σ is decreased after the first optimisation step although all grasping fingers, except for the thumb, touch the bunch of keys simultaneously (see Figure 7.32b). A rotation of the object is inhibited, but this leads to a different thumb contact point being less suitable for exerting opposing forces on the object.

Nevertheless, the cuboid model of the bunch of keys in simulation is only a coarse approximation of this form-variable object, and it was not expected that this optimisation step leads to more grasp success. On the other hand, misfortune is the reason that no grasp trial was successful, and it cannot be stated whether or not the *new pre-grasp* is disadvantageous for grasping with the real hand.

- *new target grasp:* The optimised thumb closure trajectory is slightly more directed towards the objects centre of mass in simulation. But because of the incalculable configuration of the keys, only by chance the bunch of keys was grasped more often during the evaluation of the optimised grasp with the real hand.

21. pencil:

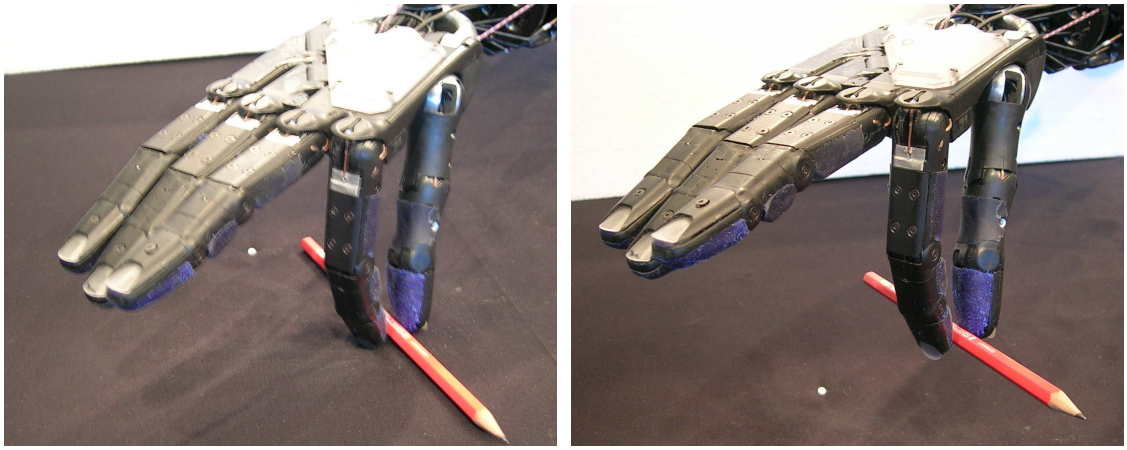

(a) *new target grasp:* the pencil is grasped. (b) *new target grasp:* the pencil is lifted.

Figure 7.33: Shadow Hand grasping the pencil in reality.

- *standard grasp:* In the experiment of grasping the pencil with the standard two finger precision grasp, the object was slightly rotated during the grasp trials. This revealed the fact that the object was touched, and additionally this standard grasp seemed to be quite appropriate for grasping the pencil. Nevertheless, none of the ten grasp trials was successful.
- *new pre-grasp:* Like in case of most other objects which are grasped with a two finger precision grasp, the first optimisation step does not lead to essential changes in the pre-grasp and also in the stability ($\sigma = 0$). Additionally, the pencil was not grasped successful in any trial.
- *new target grasp:* The second optimisation step led to the first successful trials of grasping the pencil (see Figures 7.33a and 7.33b). Additionally, the optimised thumb trajectory resulted in the largest increase in the grasp success rate over all objects. Eight out of ten grasp trials were successful in the experiment with the real hand. To find the optimal thumb target posture, frictional coefficients of $\mu = 5.0$ and $\nu = 1.33$ were used. Applying these values results in the stability σ listed in Table 7.2 ($\sigma = 0$ when using the default values presented in Section 6.1.2).

7.2 Results of the Optimisation Strategy

The experiment on human grasping described in Chapter 5 showed that humans strive for finger contact simultaneity when grasping objects. This finding was considered when developing the optimisation strategy for improving grasping with the two robot hands. The effects of the optimisation steps on the grasp times are explored in the following before the results of the experiments with the real robot hands are summarised by evaluating the first and the second optimisation step.

7.2.1 Effects on Grasp Times and Finger Contact Strategies

As proposed in Section 5.1, the grasp times to be investigated are the grasp forming time (GFT), the grasp opposing time (GOT), and the grasp closure time (GCT). Whereas GFT always equals GOT when a two finger grasp type is applied, in general the relations between the grasp times are:

$$GCT \geq GFT \geq GOT\,.$$

One result of the experiment on human grasping presented in Section 5.6 is that the most reliable measures to analyse and compare simultaneity of finger contacts when grasping objects are the ratios of GFT to GCT, and GOT to GCT:

$$GFT\% = \frac{GFT}{GCT}\,, \qquad GOT\% = \frac{GOT}{GCT}\,.$$

The grasp times and the measures of contact simultaneity (GFT% and GOT%) are determined by counting the discrete time steps during the simulated grasp closure phase.

In human grasping, the grasp closure phase starts during the reaching phase of the transport component. The hand opening in pre-grasp posture does not influence the grasp closure time (GCT). In contrast, the grasp closure starts with the hand in grasp position in simulation. Thus, a smaller hand opening results in a shorter GCT when grasping the same object, while GFT and GOT are unaffected. To take this effect into account, the changes in grasp times are discussed below besides analysing GFT% and GOT% for contact simultaneity.

Tables 7.3 and 7.4 show the values of the ratios GFT% and GOT% before optimisation and after each optimisation step. Additionally, the finger contact strategy is added, distinguishing “drag” (finger drag) from “push” (thumb push) and “both” strategies (see Section 5.1). For determination of the finger contact strategies, the *both* strategy is not only assumed to be applied when GOT% equals zero but also when the opposing forces occur in two consecutive simulation time steps. This is necessary for considering dynamical simulation effects like interpenetration as described in Section 6.1.1.

Because GFT, GCT, and GOT are not listed in Tables 7.3 and 7.4, differences in these times between the grasps of different optimisation steps are presented in the following, before the effects on grasp simultaneity are analysed.

Grasping with the TUM Hand

Differences in GCT, GFT, and GOT in the case of using the ***new pre-grasp*** compared to the *standard grasp* (compare Table 7.3) are:

- Objects no. 1, 5, 7, 20: GCT and GFT are shorter because the distance between the middle finger and the object is nearer in the *new pre-grasp* posture and both fingers now touch the object simultaneously. Although GOT% is larger, GOT is equal.

no.	object	standard grasp			new pre-grasp			new target grasp		
		type	GFT%	GOT%	type	GFT%	GOT%	type	GFT%	GOT%
1	adhesive tape	drag	50.0	23.1	drag	38.1	28.6	drag	38.1	33.3
2	toy propeller	drag	7.9	5.3	both	5.4	2.7	drag	32.7	13.5
3	toy cube	both	0.0	0.0	both	0.0	0.0	–	–	–
4	can	drag	66.7	42.9	drag	50.0	42.9	drag	53.3	40.0
5	tissue pack	drag	54.2	33.3	drag	42.1	36.8	drag	42.1	36.8
6	tennis ball	drag	39.1	34.8	drag	33.3	28.6	drag	26.3	21.1
7	paper ball	drag	54.5	31.8	drag	37.5	37.5	drag	41.2	23.5
8	sharpener	drag	–	–	drag	46.7	25.0	drag	33.3	29.2
9	remote control	drag	58.6	37.9	drag	47.8	43.5	drag	47.8	39.1
10	cup	drag	88.9	77.8	both	0.0	0.0	drag	90.9	63.6
11	board marker	drag	18.9	18.9	drag	18.9	18.9	drag	18.9	18.9
12	tea light	drag	38.3	26.7	drag	19.6	19.6	drag	27.5	25.5
13	golf ball	drag	42.9	40.0	drag	37.5	34.4	drag	35.5	32.3
14	matchbox	drag	31.8	31.8	drag	33.3	33.3	drag	44.4	44.4
15	light bulb	drag	37.9	37.9	drag	37.9	37.9	drag	41.9	41.9
16	chocolate bar	drag	39.2	29.4	drag	32.6	26.1	drag	35.4	33.3
17	folding rule	drag	21.1	21.1	drag	14.3	14.3	drag	21.1	21.1
18	voltage tester	drag	18.4	18.4	drag	20.5	20.5	drag	24.4	24.4
19	eraser	drag	21.4	21.4	drag	21.4	21.4	both	2.9	2.9
20	bunch of keys	drag	34.6	28.8	drag	30.6	30.6	drag	39.3	26.8
21	pencil	drag	–	–	drag	–	–	drag	28.6	28.6
all objects		mean: SD:	38.1 ±21.4	29.5 ±16.1	mean: SD:	31.4 ±20.4	28.9 ±15.9	mean: SD:	35.1 ±19.8	29.7 ±15.3

Table 7.3: Evaluation of TUM Hand grasp times in simulation after different steps of the optimisation strategy. No values can be determined ("–") if the grasping fingers do not touch the object simultaneously at any time. In the case of the toy cube (object no. 3), the second optimisation step is omitted because the thumb is not involved in the grasp. Red values indicate *new pre-grasps* that are used in the second optimisation step.

- Objects no. 2, 6, 12, 13, 16, 17: GCT, GFT, and GOT are shorter.
- Object no. 3: The *new pre-grasp* posture does not differ from the standard one.
- Object no. 4: GCT, GFT, and GOT are shorter, and GOT% is equal.
- Object no. 9: GCT, GFT, and GOT are shorter, but GOT% is larger because GCT decreases to a higher percentage than GOT.
- Object no. 10: GFT and GOT are zero, but the *new pre-grasp* is not applicable with the real hand (see Section 7.1.1).
- Objects no. 11, 15, 19: GCT, GFT, and GOT do not change.
- Objects no. 14, 18: GCT, GFT, and GOT are slightly longer.

In general, the first optimisation step leads to a narrower hand opening shortening the time that the fingers require for touching the object. The result is a shorter GCT, and although in most cases GFT and GOT are shorter or do not change, the division by GCT leads to small changes in GFT% and GOT%. Hence, the means of GFT% and GOT% over all objects can only be levelled down slightly with the first optimisation step.

Nevertheless, the low dexterity of the thumb is the reason why simultaneous grasping with the TUM Hand utilising the optimised pre-grasp can only be achieved in three cases. The thumb of the TUM Hand possesses two less joints when comparing with the Shadow Hand, and its closure

movement has to be directed more parallel to instead of towards the object. Thus, only when grasping the toy propeller or the cup, a *both* contact strategy can be achieved. The third grasp in which the opposing forces occur simultaneously is applied when grasping the toy cube. This case is an exceptional form of the *both* strategy because the thumb is not involved. The remaining objects have to be dragged by the other fingers before the thumb is able to touch them.

Differences in GCT, GFT, and GOT in the case of using the ***new target grasp*** compared to the *standard grasp* or to the *new pre-grasp* (if utilised, red values in Table 7.3) are:

- Object no. 1: GCT and GFT do not change. The thumb touches the object slightly later, and thus the GOT is slightly longer.
- Objects no. 2, 12, 14, 16, 18: GCT, GFT, and GOT are longer.
- Objects no. 4, 7: GCT and GOT are slightly shorter, GFT is equal, thus GFT% is slightly larger.
- Objects no. 5, 11, 15, 17: GCT, GFT, and GOT do not change.
- Objects no. 6, 13, 19: GCT, GFT, and GOT are shorter due to a thumb motion more directed towards the object.
- Object no. 8: Although GOT% is larger, GCT, GFT, and GOT are shorter due to a thumb motion more directed towards the object.
- Object no. 9: GCT and GFT do not change. The thumb touches the object slightly earlier, and thus GOT is slightly shorter.
- Objects no. 10, 20: GFT and GCT are longer, GOT is equal, thus GOT% is smaller.

When optimising the thumb opposition in the second optimisation step, grasp simultaneity is not considered. Therefore, some optimised grasps are less simultaneous, and two grasps of type *both* strategy change into the *finger drag* strategy. But in some other cases, the optimised target grasp posture shortens the distance between thumb and object, and this leads to shorter grasp times. In the case of grasping the eraser, this distance was shortened to zero, changing the former *finger drag* into a *both* strategy. Since no object is touched by the thumb before any other finger touches it, the *thumb push* contact strategy was not observed when grasping with the TUM Hand.

The large values of GFT% and GOT% in the case of large sized objects, like the can or the cup, are due to the early touch of the first finger leading to less difference of GCT compared with GFT and GOT.

Grasping with the Shadow Hand

Differences in GCT, GFT, and GOT in the case of using the ***new pre-grasp*** compared to the *standard grasp* (compare Table 7.4) are:

- Objects no. 1, 12, 14: An optimal simultaneity of finger contacts was achieved since GFT and GOT are zero.
- Objects no. 2, 4, 5, 6, 7, 8, 16: Although GFT and GOT are not equal to zero, the *new pre-grasp* realises a ”both” grasp because opposing forces occur in two consecutive simulation steps (even in the case of the object no. 6).

no.	object	standard grasp			new pre-grasp			new target grasp		
		type	GFT%	GOT%	type	GFT%	GOT%	type	GFT%	GOT%
1	adhesive tape	drag	22.2	11.1	both	0.0	0.0	drag	8.7	8.7
2	toy propeller	push	17.5	15.0	both	5.7	5.7	both	5.6	5.6
3	toy cube	push	44.7	44.7	push	44.7	44.7	push	43.7	43.7
4	can	drag	61.5	30.8	both	16.7	8.3	drag	54.5	40.9
5	tissue pack	drag	36.0	16.0	both	11.1	5.6	drag	11.1	11.1
6	tennis ball	drag	42.9	14.3	both	11.1	11.1	both	11.1	11.1
7	paper ball	drag	42.9	28.6	both	14.3	7.1	both	14.3	7.1
8	sharpener	both	16.2	2.7	both	3.1	3.1	both	3.1	3.1
9	remote control	drag	62.5	25.0	drag	21.1	15.8	drag	21.1	10.5
10	cup	drag	61.5	61.5	both	0.0	0.0	drag	61.5	61.5
11	board marker	drag	8.8	8.8	drag	8.8	8.8	drag	11.4	11.4
12	tea light	push	58.2	30.9	both	0.0	0.0	push	50.0	42.3
13	golf ball	drag	16.7	6.7	drag	13.8	6.9	drag	10.7	7.1
14	matchbox	push	48.9	35.6	both	0.0	0.0	both	0.0	0.0
15	light bulb	drag	47.2	19.4	drag	34.5	10.3	drag	20.8	12.5
16	chocolate bar	drag	26.1	6.5	both	10.5	2.6	both	15.0	0.0
17	folding rule	drag	8.8	8.8	drag	8.8	8.8	push	53.8	53.8
18	voltage tester	drag	11.4	11.4	drag	11.4	11.4	drag	18.4	18.4
19	eraser	drag	11.1	11.1	drag	11.1	11.1	push	44.4	44.4
20	bunch of keys	drag	18.2	6.8	drag	7.7	5.1	drag	20.0	17.8
21	pencil	drag	–	–	drag	–	–	drag	22.7	22.7
all objects		mean: SD:	33.2 ±19.5	19.8 ±15.0	mean: SD:	10.4 ±9.4	8.6 ±7.9	mean: SD:	30.8 ±18.0	24.7 ±15.1

Table 7.4: Evaluation of Shadow Hand grasp times in simulation after different steps of the optimisation strategy. No values can be determined ("–") when the grasping fingers do not touch the object simultaneously at any time. Red values indicate *new pre-grasps* that are used in the second optimisation step.

- Object no. 3: Since the new pre-grasp posture does not differ from the standard one, the same values for GFT and GOT are assumed.
- Objects no. 9, 15, 20: GCT, GFT and GOT are shorter.
- Object no. 10: GFT and GOT are zero, but the *new pre-grasp* is not applicable with the real hand (see Section 7.1.2).
- Objects no. 11, 17, 18, 19: GCT, GFT and GOT do not change.
- Object no. 13: GCT and GFT are shorter, GOT is equal, thus GOT% is larger.

The dexterity of the thumb of the Shadow Hand is much more improved than that of the TUM Hand. Hence, the thumb closure movement can be directed more towards the object. The benefit is that with the first optimisation step more simultaneous (*both*) grasps are realised. Although not directly comparable, the values of GFT% and GOT% over all objects after optimising for contact simultaneity are similar to those found in the experiment on human grasping (see Section 5.5). Thus, we assume that humanlike contact simultaneity is achieved in grasping with the anthropomorphic Shadow Hand when using the optimised pre-grasps.

Differences in GCT, GFT, and GOT in the case of using the ***new target grasp*** compared to the *standard grasp*, or the *new pre-grasp* (if utilised, red values in Table 7.4) are:

- Objects no. 1, 4, 5: GFT and GOT are longer due to a later contact of the thumb with the object which returns the "both" strategy into a finger drag contact strategy.

- Object no. 2: GCT is longer, GFT and GOT do not change, thus GOT% and GFT% are smaller.
- Object no. 3: GCT, GFT, and GOT are longer, although GFT% and GOT% are slightly smaller.
- Objects no. 6, 7, 8, 10, 14: GCT, GFT, and GOT do not change.
- Object no. 9: GCT and GFT do not change, and GOT is shorter.
- Objects no. 11, 12, 17, 20: GCT, GFT, and GOT are longer.
- Objects no. 13, 15: GCT and GFT are shorter, GOT is equal, thus GOT% is larger.
- Object no. 16: GCT and GFT are longer, GOT equals zero, thus GOT% equals zero too.
- Objects no. 17, 19: GCT is shorter, GOT and GFT are longer.

In general, the grasp times become longer again by optimising the thumb opposition. But although three optimised grasps change into a *thumb push* contact strategy, six grasps still realise the *both* strategy. Even when a *thumb push* or a *finger drag* strategy is used, the optimised thumb opposition leads to more stability and more successful grasp trials in most cases. But the mean values of GFT% and GOT% over all objects are similar to those of the standard grasps and do not expose the described optimisation effects.

7.2.2 First Optimisation Step

In each case in which the grasp stability value was increased by utilising the new pre-grasp, the grasp evaluation with the real hand led to more successful grasp trials or full grasp success was achieved, too. The only exception is the grasp of the cup with any of the robot hands used because the new pre-grasp is not applicable (see below). In the cases when the first optimisation step led to a lower stability value, the grasp evaluation with the real hand also resulted in a lower grasp success except when grasping the eraser. In this case, the new pre-grasps of both hands are similar to the standard pre-grasps, and also the stability values only differ slightly. Thus, the greater success in grasping with the real hands were achieved by chance.

As described in Section 6.3, the purpose of the first simulation step is, to realise finger contact simultaneity when grasping an object. If the thumb is able to reach a target object that is fixated in the simulated world, this goal can be achieved. If the thumb is not able to reach the object, at least all remaining fingers touch the object simultaneously when applying new pre-grasp. The target object is less rotated while being dragged by the fingers towards the thumb. Especially, roundish objects like the golf ball, the tennis ball, or the tea light are prevented from being pushed out of the hand because the shape of the new pre-grasp posture is adapted to the objects shape. Additionally, in most cases, the object is moved less far before being stopped by the thumb. This is due to a more flexed thumb pre-grasp posture so that the thumb is nearer its target grasp posture when it touches the object. Less rotational and translational motion of the object while being grasped leads to larger grasp stability σ in general.

The larger flexion of the fingers in the new pre-grasp posture also leads to a narrower hand opening enlarging the constraints on variability in position and orientation of the real robot hand. These constraints have been met when evaluating the optimised grasps with the real robot hands except in the case of grasping the cup. Even if optimal positioning of the cup would be achieved, the new pre-grasp posture of the approaching robot hand is that narrow that at least one finger would

stick into the cup (for detailed explanation, see Sections 7.1.1 and 7.1.2). The problem is the large size and the specifically conical shape of the cup being smaller at its bottom than at its top. But although the new pre-grasp posture was not applicable in the case of the cup, it was shown that the grasp can be improved by the second optimisation step leading to full grasp success in the case of both real robot hands.

7.2.3 Second Optimisation Step

Optimising the thumb closure trajectory in the second optimisation step always led to a larger grasp stability σ. With a thumb in optimised opposition posture, the stability can increase by a multiple when comparing to the stability value of the standard grasp or that of the grasp optimised in the first step. The reason is that the thumb provides the major part of the opposing force to each force exerted by the other fingers.

The evaluation of the new target grasps applied with the real hands led to a larger number of successful grasps or even to full success in each case. In the only exception, resulting in nine successful grasp trials compared to ten after the first optimisation step, the reasons for the incomplete grasp trial are the variances in object position and orientation and the inaccuracies in the finger control of the TUM Hand. In general, we observed that although contact simultaneity is not considered when optimising the target grasp, the optimised thumb target grasp posture leads to less rotation of the object after being touched by the thumb during grasp closure.

In general, the optimised thumb closure trajectory leads to a contact point near the centre of mass of the target object. In the case of grasping a flat object with a less dextrous hand, like the matchbox being grasped by the TUM Hand (see Section 7.1.1), this optimal contact point cannot be reached. By applying the evolutionary algorithm in the second optimisation step, however, a thumb target posture is learned that leads to the most optimal and reachable contact point which, in the case of the matchbox, is farther from and lower than its centre of mass. This optimal result is hardly achievable by traditionally approaches in which grasp points have to be determined in a computationally demanding way and have to be matched by the fingertips via calculation of inverse kinematics.

7.3 Comparison of TUM and Shadow Hand

Comparing the numbers of successful grasp trials of the TUM and the Shadow Hand before optimisation (row ”standard grasp” in Table 7.5) shows that the Shadow Hand is able to grasp more objects with higher reliability. This advantageous suitability for grasping is not expressed in the success values after optimisation listed in Tables 7.1 and 7.2 except in the cases of the bunch of keys and the pencil. However, these values were determined by using fingertip force sensors in the case of the TUM Hand and without any sensors for touch detection in the case of the Shadow Hand.

To compare both hands in terms of capability for grasping and to judge the influence of touch sensors, we performed an additional experiment for evaluation of the optimised grasps with the TUM Hand. The fingertip sensors of the TUM Hand were not utilised in this experiment. In this case, the finger controller proceeds in trying to reach the target grasp posture even when the fingers already touch the object. The resulting numbers of successful grasp trials are listed as second value in the second last column of Table 7.5.

no.	object	grasp type	standard grasp		optimised	
			TUM	Shadow	TUM	Shadow
1	adhesive tape	power t_3	10	10	10 / 9	10
2	toy propeller	3F spec t_5	10	10	10 / 10	10
3	toy cube	2F pinch t_4	10	10	10 / 0	10
4	can	power t_3	10	10	10 / 2	10
5	tissue pack	power t_3	10	10	10 / 10	10
6	tennis ball	power t_3	10	7	10 / 4	10
7	paper ball	power t_3	9	10	10 / 10	10
8	sharpener	AF prec t_1	8	10	10 / 10	10
9	remote control	power t_3	8	10	10 / 10	10
10	cup	power t_3	9	10	10 / 0	10
11	board marker	2F prec t_2	7	10	10 / 10	10
12	tea light	AF prec t_1	6	8	10 / 3	10
13	golf ball	power t_3	7	6	10 / 6	9
14	matchbox	AF prec t_1	7	6	9 / 4	10
15	light bulb	power t_3	6	8	10 / 0	10
16	chocolate bar	AF prec t_1	5	10	10 / 10	10
17	folding rule	2F prec t_2	4	10	10 / 10	10
18	voltage tester	2F prec t_2	3	8	9 / 8	9
19	eraser	2F prec t_2	4	9	10 / 10	10
20	bunch of keys	AF prec t_1	0	1	0 / 0	2
21	pencil	2F prec t_2	0	0	0 / 0	8

Table 7.5: Number of successful grasp trials (0 to 10) out of 10 attempts with both real hands before and after optimisation. The second value, in the case of the optimised TUM grasps, denote the results of the experiment in which the fingertip sensors were not utilised.

Results of the Additional Experiment

In this additional experiment, most objects were grasped with nearly full success, but in eight cases the grasp success decreased significantly. One object that cannot be grasped without touch detection is the toy cube. This object was tried to be grasped with the index finger and the middle finger (two finger pinch t_4). The combined control of the first two finger joints leads to a redirection of the adduction movement into a flexion or extension movement when both fingers touch the toy cube. The result is the loss of the object if the fingers are not stopped. The seven remaining cases can be divided into two groups. The first group consists of all large benchmark objects which are touched by other phalanges besides the fingertips during grasp closure (objects no. 4, 6, 10, and 15). Because the index and the middle finger exert forces with their proximal or middle phalanges that the tip of the thumb can hardly resist, in some grasp trials the target object was pushed out of the hand. All small objects being grasped by more than two fingers (objects no. 12, 13, and 14) compose the second group of objects that are difficult to be grasped with the TUM Hand. These objects have to be optimally positioned under the grasping TUM hand if no force sensors are used. But because of variances in object position and orientation and the inaccuracies in the finger control of the TUM Hand, these objects are lost in some cases.

Advantageous of the Shadow Hand

The results of the additional experiment lead to the advantages of the Shadow Hand compared to the TUM Hand. Although no touch sensors were used in the case of the Shadow Hand too, most objects were grasped with more success when grasping with it. This fact can be explained by two main reasons. The first one is the dexterity and the shape of the thumb. Because the shape of the

Shadow thumb is much more anthropomorphic while possessing two more joints compared to the thumb of the TUM Hand, it can be optimally opposed to the other fingers. Opposing forces can be better resisted not only by the tip of the thumb but also by its middle phalanx. Secondly, the Shadow Hand has two more fingers which realise a cage-like enclosure of the object in a power grasp or an all finger precision grasp. Hence, larger variances in object position and orientation relative to the Shadow Hand being in grasp position can be tolerated when grasping small objects. Additionally, more grasping fingers lead to more contact points and thus more forces with different directions are exerted on large objects.

8 Conclusions

This thesis presented a complete line of biologically motivated approaches for providing robot hands with grasping capabilities. These approaches comprise a grasp synthesis, a grasp strategy, and a grasp taxonomy.

The grasp types defined by the taxonomy are fairly easy to realise in a robot hand setup when following the development rules we propose. Our approach to grasp synthesis stresses the target grasp posture providing the opportunity for optimising the realised grasp types for finger closure trajectories. Besides optimising the target grasp, the optimisation strategy we propose optimises the pre-grasp posture for contact simultaneity and is substantiated by an experiment on human grasping.

By implementing the grasp strategy and the optimisation strategy on one robot hand setup and porting these strategies into the second setup, we showed that these strategies are realisable on and portable among totally different robot systems including anthropomorphic robot hands with different numbers of fingers. The strategies proposed are robust against limited positioning accuracy of the finger joints and uncertainties about object position and orientation.

In this final chapter, we summarise conclusions about the optimisation strategy and provide an outlook on potential improvements.

Optimisation Strategy

Optimising in simulation leads to additional inaccuracies due to limitations of the simulator. Nevertheless, the evaluation of the optimisation strategy showed that the pre-grasp postures and target grasp postures optimised in simulation result in higher grasp stability and in more successful grasp trials when applied with the real hands. Except for objects that are difficult to grasp (bunch of keys and pencil), nearly all trials of grasping a benchmark object with any of both robot hands used was successful when applying the combined results of both optimisation steps.

The first optimisation step results in larger stability values and more grasp success in many cases. But optimising a two finger precision grasp did not lead to contact simultaneity in any case. The reason is that the thumb closure trajectory ends before touching the object. The index finger has to drag the target object towards the thumb and this fact cannot be influenced by the first optimisation step. In the case of other grasp types, more fingers are involved and even if the thumb cannot touch the target object at the same point in time as the other fingers, at least these fingers touch the object simultaneously. This leads to a larger grasp stability in most cases. Nevertheless, it was shown that if the stability can be enhanced the new pre-grasp is suitable for grasping with the real hand.

The second optimisation step optimised the thumb trajectory and led to a larger grasp stability in any case. The thumb opposition is very important for successful grasping since the thumb is mostly responsible for resisting the forces exerted by all other fingers. The evolutionary algorithm applied for learning the thumb joint angle values has proven to be a very suitable and fast learning method for this task. The evaluation of the optimised target grasps indicated that the results of the second optimisation step determined in simulation leads to optimal grasps when applied with the real hands in every case.

Outlook

Two objects of the benchmark system we propose are hard to be grasped with the robot hands used. One of these objects is the pencil. After optimising the standard grasp for grasping this object, at least the dextrous Shadow Hand was able to grasp the pencil successfully. But in the case of the TUM Hand, successful grasping was only possible after elevating the object 4mm above the desktop surface. The reason is a remaining gap between the roundish tips of the thumb and the index finger in target grasp posture. To close this gap a grasp formed like a straight-tip tweezers is required. But this implies a less distance between the palm and the object and thus a different position of the hand. We suppose that a suitable grasp can be realised by optimising the standard two finger pinch grasp for the height of the palm and the flexion of the fingers. To this end, five interdependent parameters have to tuned. For considering these parameters, the optimisation strategy is easily expandable, and we assume that the evolutionary algorithm we propose is suitable for finding an optimal solution in this five-dimensional space.

Solving the task of grasping the second object, the bunch of keys, with a robot hand is even more difficult. Only a few grasps trials with the Shadow Hand were successful. The arrangement of the keys during grasp closure is not predictable, and in most cases the grasping fingers slip over the plain surface of a key. But we assume that this kind of object can be grasped successfully by realising a humanlike grasping process in which the fingertips glide over the desktop while aligning the keys. For this purpose, a real-time force control of the finger joints has to be realised preferably supported by a force control of the robot arm for adjusting the position and orientation of the palm. Furthermore, sensors for slip detection in the fingertips would be beneficial. Besides the development and the implementation of the sensors required and the realisation of force controllers, optimising the finger trajectories under these conditions leaves much room for future work.

A Joint Values

The appendix presents tables listing joint angle values of the robot hands used. Values of the standard grasps are given in Section A.1. Section A.2 presents the tables listing the values of the optimised grasps and the gene bounds used for the second optimisation step.

A.1 Standard Grasps

grasp type	thumb			index finger			middle finger		
t	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	θ_9
pre-grasp									
all finger precision t_1	0	0	0	0	0	0	0	0	0
two finger precision t_2	4	12	18	0	12	20	0	0	0
power t_3	10	16	24	0	0	19	0	0	19
two finger pinch t_4	0	0	0	8	3	0	-8	3	0
three finger special t_5	0	0	0	8	3	0	-8	3	0
target grasp									
all finger precision t_1	2	42	4	-5	46	22	5	48	23
two finger precision t_2	6	38	0	0	46	25	0	0	0
power t_3	10	28	24	0	30	43	-2	32	44
two finger pinch t_4	0	0	0	-7	3	0	7	3	0
three finger special t_5	0	46	0	2	30	32	5	30	32

Table A.1: Joint angle values of the TUM Hand standard grasps in degrees.

grasp type	thumb					index finger			middle finger			ring finger			little finger			palm
	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	θ_9	θ_{10}	θ_{11}	θ_{12}	θ_{13}	θ_{14}	θ_{15}	θ_{16}	θ_{17}	θ_{18}
pre-grasp																		
all finger precision t_1	10	53	0	8	25	0	0	19	0	0	19	0	0	19	0	0	19	0
two finger precision t_2	-5	55	0	10	0	0	12	20	0	0	0	0	0	0	0	0	0	0
power t_3	10	54	0	8	45	0	0	19	0	0	19	0	0	19	0	0	19	0
two finger pinch t_4	16	15	0	5	5	-2	65	75	0	0	0	0	0	0	0	0	0	0
three finger special t_5	0	0	0	0	0	20	3	0	-20	3	0	-17	0	0	-15	0	0	0
target grasp																		
all finger precision t_1	-23	73	0	16	-10	20	81	12	6	81	13	-12	82	13	-20	82	14	5
two finger precision t_2	-21	63	0	10	-9	0	65	20	0	0	0	0	0	0	0	0	0	0
power t_3	-25	77	0	16	20	0	55	40	0	55	40	0	55	40	0	55	40	0
two finger pinch t_4	15	19	0	28	26	2	65	75	0	0	0	0	0	0	0	0	0	0
three finger special t_5	-20	60	0	12	-7	20	65	20	-20	65	20	-17	30	32	-15	30	32	0

Table A.2: Joint angle values of the Shadow Hand standard grasps in degrees.

A.2 Optimised Grasps

gene	AF prec t_1		2F prec t_2		power t_3		2F pinch t_4		3F spec t_5	
bounds	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2
TUM Hand										
lower	-4	20	-10	15	-12	6	–	–	-10	15
upper	12	50	12	45	12	35	–	–	12	50
Shadow Hand										
lower	-30	55	-30	50	-30	50	-15	10	-35	45
upper	0	85	-15	70	0	85	35	40	-5	85

Table A.3: Gene bounds ($i_1 = \theta_1$; $i_2 = \theta_2$) for optimising the target grasps of the TUM Hand and the Shadow Hand with the evolutionary algorithm. Values are given in degrees.

object	grasp	after optimisation step **one**									**two**	
		thumb			index finger			middle finger			thumb	
	type	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	θ_9	θ_1	θ_2
adhesive tape	t_3	10.0	18.7	24.0	0.0	0.0	19.0	-0.3	3.4	23.0	11.2	28.1
toy propeller	t_5	0.0	12.7	0.0	8.0	3.0	0.0	-7.4	4.1	2.4	8.5	39.7
toy cube	t_4	0.0	0.0	0.0	8.0	3.0	0.0	-8.0	3.0	0.0	–	–
can	t_3	10.0	20.6	24.0	0.0	0.0	19.0	-0.4	5.4	26.4	-10.7	15.2
tissue pack	t_3	10.0	19.9	24.0	0.0	0.0	19.0	-0.3	3.4	23.9	9.3	24.8
tennis ball	t_3	10.0	18.6	24.0	0.0	0.0	19.0	-0.3	2.4	22.3	6.9	29.0
paper ball	t_3	10.0	19.4	24.0	0.0	0.0	19.0	-0.2	3.4	23.9	9.0	28.3
sharpener	t_1	0.4	9.2	0.0	0.0	0.0	0.0	0.4	4.1	0.0	-1.3	41.0
remote control	t_3	10.0	20.5	24.0	0.0	0.0	19.0	-0.3	4.1	24.8	10.0	25.0
cup	t_3	10.0	23.4	24.0	0.0	0.0	19.0	-0.4	3.0	23.3	11.1	29.4
board marker	t_2	4.5	16.4	18.0	0.0	12.0	20.0	0.0	0.0	0.0	6.0	33.0
tea light	t_1	0.4	7.9	0.0	0.0	0.0	0.0	0.3	3.6	0.0	5	36.4
golf ball	t_3	10.0	16.0	24.0	0.0	0.0	19.0	-0.4	3.8	24.1	9.3	32.7
matchbox	t_1	0.5	9.2	0.0	0.0	0.0	0.0	0.5	4.8	0.0	8.6	35.1
light bulb	t_3	10.0	17.1	24.0	0.0	0.0	19.0	0.0	0.0	19.0	11.1	25.3
chocolate bar	t_1	0.5	9.2	0.0	0.0	0.0	0.0	0.6	5.4	0.0	5.2	39.9
folding rule	t_2	4.5	16.4	18.0	0.0	12.0	20.0	0.0	0.0	0.0	6.5	31.1
voltage tester	t_2	4.5	17.1	18.0	0.0	12.0	20.0	0.0	0.0	0.0	7.1	30.1
eraser	t_2	4.4	17.1	18.0	0.0	12.0	20.0	0.0	0.0	0.0	-5.5	40.1
bunch of keys	t_1	0.5	9.8	0.0	0.0	0.0	0.0	0.6	5.4	0.0	1.9	41.5
pencil	t_2	5.7	33.7	0.0	0.0	12.0	20.0	0.0	0.0	0.0	3.7	30.6

Table A.4: Joint angle values of the TUM Hand optimised for grasping the benchmark objects. All joint angles of the pre-grasp are optimised in optimisation step **one**. The first two thumb angles of the target grasp are optimised in optimisation step **two**.

object	grasp type	after optimisation step **one**																		**two**	
		thumb					index finger			middle finger			ring finger			little finger			palm	thumb	
		θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	θ_9	θ_{10}	θ_{11}	θ_{12}	θ_{13}	θ_{14}	θ_{15}	θ_{16}	θ_{17}	θ_{18}	θ_1	θ_2
adhesive tape	t_3	3.7	58.3	0.0	8.3	43.5	0.1	3.9	19.7	0.0	0.0	19.0	0.0	2.3	19.4	0.0	10.0	21.0	0.0	-19.1	73.3
toy propeller	t_5	0.0	0.0	0.0	0.0	0.0	20.0	6.8	0.7	-20.0	9.7	1.3	-17.0	30.0	9.5	-15.0	30.0	9.5	0.0	-20.1	59.4
toy cube	t_4	16.0	15.0	0.0	5.0	5.0	-2.0	65.0	75.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	19.6	20.1
can	t_3	-9.7	67.2	0.0	9.0	40.9	0.0	0.0	19.0	0.0	6.6	20.4	0.0	8.0	20.6	0.0	9.8	20.9	0.0	-14.9	79.1
tissue pack	t_3	-2.5	62.4	0.0	8.6	42.3	0.0	0.0	19.0	0.0	6.5	20.1	0.0	7.9	20.4	0.0	8.1	20.4	0.0	-24.3	78.9
tennis ball	t_3	0.3	60.6	0.0	8.4	42.9	0.2	8.2	20.6	0.0	0.0	19.0	0.0	0.0	19.0	0.0	12.9	21.6	0.0	-27.2	78.1
paper ball	t_3	-4.0	63.4	0.0	8.7	42.0	0.0	0.0	19.0	0.0	3.7	19.8	0.0	6.5	20.3	0.0	6.9	20.4	0.0	-23.9	74.8
sharpener	t_1	10.1	53.0	0.0	8.0	24.7	3.7	14.2	19.0	0.0	0.0	19.0	0.9	5.2	19.0	-10.5	43.2	14.0	2.9	-20.2	67.3
remote control	t_3	3.0	58.6	0.0	8.4	43.5	0.0	0.0	19.0	0.0	8.0	20.5	0.0	10.9	21.0	0.0	12.5	21.3	0.0	-25.8	73.5
cup	t_3	-8.2	66.3	0.0	9.0	41.2	0.0	3.7	19.8	0.0	0.0	19.0	0.0	0.0	19.0	0.0	7.1	20.4	0.0	-15.0	67.8
board marker	t_2	-5.0	55.0	0.0	10.0	-0.4	0.0	12.0	20.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-19.2	62.0
tea light	t_1	10.0	53.0	0.0	8.0	25.0	8.3	32.8	18.2	1.7	22.5	18.8	-4.6	31.7	13.0	-7.7	31.7	14.0	2.2	-24.8	76.8
golf ball	t_3	9.2	54.6	0.0	8.2	44.4	0.0	14.3	21.8	0.0	0.0	19.0	0.0	0.0	19.0	0.0	20.5	27.0	0.0	-27.0	77.5
matchbox	t_1	10.0	53.0	0.0	8.0	25.0	5.0	20.0	18.2	1.7	22.6	18.9	-3.8	25.4	18.8	-12.3	50.3	19.0	2.8	-25.1	74.3
light bulb	t_3	3.0	58.8	0.0	8.4	43.5	0.0	5.1	19.9	0.0	8.0	20.5	0.0	3.7	19.7	0.0	0.0	19.0	0.0	-28.1	73.8
chocolate bar	t_1	10.1	53.1	0.0	8.0	24.7	0.0	0.0	19.0	0.6	8.0	18.8	-1.7	10.8	18.8	-1.7	6.7	18.8	0.0	-22.9	76.0
folding rule	t_2	-4.8	55.1	0.0	9.9	-0.4	0.0	12.0	20.0	0.0	0.0	0.0	0.0	0.	0.0	0.0	0.0	0.0	0.0	-30.0	51.7
voltage tester	t_2	-4.8	55.1	0.0	9.9	-0.2	0.0	12.0	20.0	0.0	0.0	0.0	0.0	0.	0.0	0.0	0.0	0.0	0.0	-19.3	63.9
eraser	t_2	-4.9	55.0	0.0	10.0	0.0	0.0	12.0	20.0	0.0	0.0	0.0	0.0	0.	0.0	0.0	0.0	0.0	0.0	-27.3	53.2
bunch of keys	t_1	10.2	53.0	0.0	8.0	24.4	0.0	0.0	19.0	0.6	8.0	18.8	-1.7	10.8	18.8	-12.2	50.4	14.0	3.4	-22.0	75.5
pencil	t_2	-5.0	55.0	0.0	10.0	0.0	0.0	12.0	20.0	0.0	0.0	0.0	0.0	0.	0.0	0.0	0.0	0.0	0.0	-23.1	61.9

Table A.5: Joint angle values of the Shadow Hand optimised for grasping the benchmark objects. All joint angles of the pre-grasp are optimised in optimisation step **one**. The first two thumb angles of the target grasp are optimised in optimisation step **two**.

Bibliography

J. L. Alberts, M. Saling, and G. E. Stelmach. Alterations in transport path differentially affect temporal and spatial movement parameters. *Experimental Brain Research*, 143(4):417–425, April 2002.

M. A. Arbib. Perceptual structures and distributed motor control. In V.B. Brooks, editor, *Handbook of Physiology*, volume II of *Section 2: The Nervous System, Motor Control, Part 1*, pages 1449–1480. American Physiological Society, 1981.

M. A. Arbib, T. Iberall, and D. Lyons. Coordinated control programs for control of the hands. In A. W. Goodwin and I. Darian-Smith, editors, *Hand function and the neocortex*, Experimental Brain Research Supplemental 10, pages 111–129. Springer-Verlag, Berlin, 1985.

G. A. Bekey, R. Tomovic, and I. Zeljkovic. Control architecture for the belgrade/usc hand. In S. T. Venkataraman and T. Iberall, editors, *Dexterous Robot Hands*, pages 136–149. Springer-Verlag, Berlin, 1990.

N. Bennis and A. Roby-Brami. Coupling between reaching movement direction and hand orientation for grasping. *Brain Research*, 952(2):257–267, October 2002.

A. Bicchi and V. Kumar. Robotic grasping and contact: A review. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 348–353, San Francisco, 2000. IEEE.

M. Biegstraaten, J. B. Smeets, and E. Brenner. The relation between force and movement when grasping an object with a precision grip. *Experimental Brain Research*, 171(3):347–357, May 2006.

R. J. Bootsma, R. G. Marteniuk, C. L. MacKenzie, and F. T. J. M. Zaal. The speed-accuracy trade-off in manual prehension: effects of movement amplitude, object size and object width on kinematic characteristics. *Experimental Brain Research*, 98(3):535–541, 1994.

C. Borst, M. Fischer, and G. Hirzinger. A fast and robust grasp planner for arbitrary 3d objects. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 1890 – 1896, Detroit, Michigan, USA, May 1999. IEEE.

C. Borst, M. Fischer, and G. Hirzinger. Calculating hand configurations for precision and pinch grasps. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 1553–1559, Lausanne, Switzerland, October 2002. IEEE.

C. Borst, M. Fischer, and G. Hirzinger. Grasping the dice by dicing the grasp. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 3692–3697, Las Vegas, Nevada, USA, October 2003. IEEE.

C. Borst, M. Fischer, and G. Hirzinger. Calculating human like grasp shapes: Pinch grasps. In E. Prassler, G. Lawitzky, A. Stopp, G. Grunwald, M. Hägele, R. Dillmann, and I. Iossifidis, editors, *Advances in Human-Robot Interaction*, volume 14 of *Tracts in Advanced Robotics*, pages 181–193. Springer, 2004.

C. Borst, M. Fischer, and G. Hirzinger. Efficient and precise grasp planning for real world objects. In F. Barbagli, D. Prattichizzo, and K. Salisbury, editors, *Multi-point Interaction with Real and Virtual Objects*, volume 18 of *Tracts in Advanced Robotics*, chapter 6, pages 91–111. Springer, 2005.

A. W. Burton and M. J. Dancisak. Grip form and graphomotor control in preschool children. *American Journal of Occupational Therapy*, 54(1)::9–17, January 2000.

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.

G. Butterworth and S. Itakura. Development of precision grips in chimpanzees. In *Developmental Science*, volume 1, pages 39–43. Blackwell Publishers Ltd., April 1998.

A. Caffaz and G. Cannata. The design and development of the dist-hand dextrous gripper. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*. IEEE, May 1998.

S. Chieffi and M. Gentilucci. Coordination between the transport and the grasp components during prehension movements. *Experimental Brain Research*, 94(3):471–477, 1993.

M. Ciocarlie, A. Miller, and P. K. Allen. Grasp analysis using deformable fingers. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 3579–3585, Edmonton, August 2005. IEEE.

CMLabs. Vortex – physics engine for real-time simulation, 2006. URL `http://www.cm-labs.com/products/vortex/`.

K. J. Cole and J. H. Abbs. Kinematic and electromyographic responses to perturbation of a rapid grasp. *Journal of Neurophysiology*, 1987.

M. R. Cutkosky. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Transactions on Robotics and Automation*, 5(3):269–279, June 1989.

M. R. Cutkosky and R. D. Howe. Human grasp choice and robotic grasp analysis. In S. T. Venkataraman and T. Iberall, editors, *Dexterous Robot Hands*, chapter 1, pages 5–31. Springer-Verlag, Berlin, 1990.

M. R. Cutkosky and P. K. Wright. Modeling manufacturing grips and correlations with the design of robotic hands. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, volume 3, pages 1533–1539. IEEE, April 1986.

M. Desmurget, C. Prablanc, M. Arzi, Y. Rossetti, Y. Paulignan, and C. Urquizar. Integrated control of hand transport and orientation during prehension movements. *Experimental Brain Research*, 110(2):265–278, July 1996.

W. Dunham. *Journey through Genius: The Great Theorems of Mathematics*, chapter Heron's Formula for Triangular Area, pages 113–132. Wiley, New York, 1990.

A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, 2003.

EURON. European robotics research network, 2006. URL `http://www.euron.org/`.

J. Fan, J. He, and S. I. H. Tillery. Control of hand orientation and arm movement during reach and grasp. *Experimental Brain Research*, pages 1–14, November 2005. Epub ahead of print.

R. Fearing. Simplified grasping and manipulation with dextrous robot hands. *IEEE Transactions on Robotics and Automation*, 2(4):188–195, December 1986.

C. Ferrari and J. Canny. Planning optimal grasps. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 2290–2295, Nice, France, May 1992. IEEE.

G. A. Fink, J. Fritsch, N. Leßman, H. Ritter, G. Sagerer, J. J. Steil, and I. Wachsmuth. *Situated Communication*, chapter Architectures of Situated Communicators: From Perception to Cognition to Learning. Mouton de Gruyter, 2006.

A. E. Flatt, MD, FRCS, and FACS. Grasp. *Baylor University Medical Center*, 13(4):343–348, October 2000.

D. B. Fogel. An introduction to simulated evolutionary optimization. *IEEE Transactions on Neural Networks*, 5(1):3–14, January 1994.

L. J. Fogel, A. J. Owens, and M. J. Walsh. *Artificial Intelligence through Simulated Evolution*. Wiley Publishing, New York, 1966.

M. Gentilucci, U. Castiello, M. L. Corradini, M. Scarpa, C. Umilta, and G. Rizzolatti. Influence of different types of grasping on the transport component of prehension movements. *Neuropsychologia*, 29(5):361–378, 1991.

M. Gentilucci, S. Chieffi, M. Scarpa, and U. Castiello. Temporal coupling between transport and grasp components during prehension movements: effects of visual perturbation. *Behavioural Brain Research*, 47(1):71–82, March 1992.

M. Gentilucci, E. Daprati, M. Gangitano, M. C. Saetti, and I. Toni. On orienting the hand to reach and grasp an object. *NeuroReport*, 7(2):589–592, January 1996.

G. Granosik and J. Borenstein. Minimizing air consumption of pneumatic actuators in mobile robots. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 3634–3639, New Orleans, 2004.

G. Granosik and E. Jezierski. Application of a maximum stiffness rule for pneumatically driven legs of walking robot. In *Proc. CLAWAR*, pages 213–218, 1999.

H. E. Griffiths. Treatment of the injured workman. *The Lancet*, 244:729–733, June 1943.

H. Hanafusa and H. Asada. Stable prehension by a robot hand with elastic fingers. In *Proc. of the Intern. Symp. on Industrial Robots*, pages 361–368, Tokyo, October 1977.

R. Haschke, J. J. Steil, I. Steuwer, and H. Ritter. Task-oriented quality measures for dextrous grasping. In *Proc. Conference on Computational Intelligence in Robotics and Automation*, Espoo, 2005. IEEE.

M. Henschel. Optimisation of grasping strategies using evolutionary algorithms. Master's thesis, Faculty of Technology, Bielefeld University, 2006.

A. Hildebrandt, O. Sawodny, R. Neumann, and A. Hartmann. A flatness based design for tracking control of pneumatic muscle actuators. In *Proc. ICARCV*, volume 3, pages 1156–1161, Singapore, 2002.

J. Holland. *Adaptation in Natural and Artificial Systems: An Inroductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, Ann Arbor, USA, 1975.

W. D. Hopkins, C. Cantalupo, M. J. Wesley, A. B. Hostetter, and D. L. Pilcher. Grip morphology and hand use in chimpanzees (pan troglodytes): Evidence of a left hemisphere specialization in motor skill. *Journal of Experimental Psychology*, 131(3):412–423, September 2002.

M. Huber and R. A. Grupen. Robust finger gaits from closed-loop controllers. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, volume 2, pages 1578–1584. IEEE, 2002.

T. Iberall. The nature of human prehension: Three dextrous hands in one. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 396–401, Raleigh, NC, 1987. IEEE.

T. Iberall. Human prehension and dexterous robot hands. *International Journal of Robotics Research*, 16(3):285 – 299, June 1997.

T. Iberall and D. M. Lyons. Towards perceptual robotics. In *International Conference on Systems, Man, and Cybernetics*, pages 147–157, Halifax, Nova Scotia, October 1984. IEEE.

T. Iberall and C. L. MacKenzie. Opposition space and human prehension. In S. T. Venkataraman and T. Iberall, editors, *Dexterous Robot Hands*, chapter 2, pages 32–54. Springer-Verlag, Berlin, 1990.

T. Iberall, G. Bingham, and M. A. Arbib. Opposition space as a structuring concept for the analysis of skilled hand movements. In H. Heuer and C. Fromm, editors, *Generation and Modulation of Action Patterns*, Exp Brain Res Series 15, pages 158–173. Springer, Berlin Heidelberg New York, 1986.

S. C. Jacobsen, E. K. Iversen, D. F. Knutti, R. T. Johnson, and K. B. Biggers. Design of the utah/mit dexterous hand. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 1520–1532. IEEE, 1986.

M. Jeannerod. Intersegmental coordination during reaching at natural visual objects. In J. Long and A. Baddeley, editors, *Attention and Performance*, volume IX, pages 153–168. Erlbaum, Hillsdale, 1981.

M. Jeannerod. The timing of natural prehension movements. *Journal of Motor Behavior*, 16(3): 235–254, 1984.

J. Jockusch. *Exploration based on Neural Networks with Applications in Manipulator Control*. PhD thesis, Faculty of Technology, Bielefeld University, 2000.

S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

Y. Kurita, J. Ueda, Y. Matsumoto, and T. Ogasawara. Cpg-based manipulation: generation of rhythmic finger gaits from human observation. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 1209–1214, New Orleans, LA, USA, April 2004. IEEE.

Y. Li and N. Pollard. A shape matching algorithm for synthesizing humanlike enveloping grasps. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids 2005)*, pages 442–449, 2005.

C. M. Light, P. H. Chappell, and B. S. Ellis. A critical review of functionality assessment in natural and prosthetic hands. *British Journal of Occupational Therapy*, 62(1):7–11, January 1999.

C. M. Light, P. H. Chappell, and P. J. Kyberd. Establishing a standardized clinical assessment tool of pathologic and prosthetic hand function: Normative data, reliability, and validity. *Archive of Physical Medicine and Rehabilitation*, 2002.

C. S. Lovchik and M. A. Diftler. The robonaut hand: A dexterous robot hand for space. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, volume 2, pages 907–912, Detroit, Michigan, May 1999. IEEE.

D. Lyons. A simple set of grasps for a dextrous hand. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 588–593. IEEE, March 1985.

C. L. MacKenzie and T. Iberall. *The grasping hand*, volume 104 of *Advances in Psychology*. North-Holland (Elsevier Science B.V.), Amsterdam, 1994.

W. J. Mallon, H. R. Brown, and J. A. Nunley. Digital ranges of motion: Normal values in young adults. *Journal of Hand Surgery*, 16A(5):882–887, 1991.

P. Mamassian. Prehension of objects oriented in three-dimensional space. *Experimental Brain Research*, 114(2):235–245, April 1997.

M. W. Marzke. Evolutionary development of the human thumb. *Hand Clinics*, 8(1):1–8, 1992.

C. R. Mason, L. S. Theverapperuma, C. M. Hendrix, and T. J. Ebner. Monkey hand postural synergies during reach-to-grasp in the absence of vision of the hand and object. *Journal of Neurophysiology*, 91(6):2826–2837, June 2004.

E. D. McBride. *Disability Evaluation*. J.B. Lippincott Co., 1942.

M. N. McDonnell, M. C. Ridding, S. C. Flavel, and T. S. Miles. Effect of human grip strategy on force control in precision tasks. *Experimental Brain Research*, 161(3):368–373, March 2005.

P. McGuire, J. Fritsch, H. Ritter, J. J. Steil, F. Röthling, G. A. Fink, S. Wachsmuth, and G. Sagerer. Multi-modal human-machine communication for instructing robot grasping tasks. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 1082–1089. IEEE, 2002.

G. A. Medrano-Cerda, C. J. Bowler, and D. G. Caldwell. Adaptive position control of antagonistic pneumatic muscle actuators. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 378–384, 1995.

R. Menzel, K. Woelfl, and F. Pfeiffer. The development of a hydraulic hand. In *Proc. 2nd Conf. Mechatronics and Robotics*, pages 225–238, 1993.

Z. Michalewicz. *Genetic algorithms + data structures = evolution programs*. Springer-Verlag, 3rd edition, 1996.

A. Miller and P. K. Allen. Graspit!: A versatile simulator for robotic grasping. *IEEE Robotics and Automation Magazine*, 11(4):110–122, December 2004.

A. Miller, P. K. Allen, V. Santos, and F. Valero-Cuevas. From robot hands to human hands: A visualization and simulation engine for grasping research. *Industrial Robot*, 32(1):55–63, 2005.

A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen. Automatic grasp planning using shape primitives. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 1824–1829, Tapei, September 2003. IEEE.

A. Morales. *Learning to predict grasp reliability with a multifinger robot hand by using visual features*. PhD thesis, Jaume-I University of Castelló, Spain, 2003.

A. Morales. Experimental benchmarking of grasp reliability. http://www.robot.uji.es/people/morales/experiments/benchmark.html, 2006.

T. Mouri, H. Kawasaki, K. Yoshikawa, J. Takai, and S. Ito. Anthropomorphic robot hand: Gifu hand iii. In *Proc. of Int. Conf. ICCAS2002*, pages 1288–1293, 2002.

M. R. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

J. R. Napier. The prehensile movements of the human hand. *Journal of Bone and Joint Surgery*, 38B(4):902–913, 1956.

K. M. Newell, P. V. McDonald, and R. Baillargeon. Body scale and infant grip configurations. *Developmental Psychobiology*, 26(4):195–205, May 1993.

V.-D. Nguyen. The synthesis of stable grasps in the plane. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, volume 3, pages 884–889. IEEE, April 1986.

V.-D. Nguyen. Constructing force-closure grasps. *International Journal of Robotics Research*, 7 (3):3–16, 1988.

A. M. Okamura, N. Smaby, and M. R. Cutkosky. An overview of dexterous manipulation. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS), Symposium on Dexterous Manipulation*, volume 1, pages 255–262. IEEE, 2000.

Y. Paulignan, M. Jeannerod, C. MacKenzie, and R. Marteniuk. Selective perturbation of visual input during prehension movements. 2. the effects of changing object size. *Experimental Brain Research*, 87(2):407–420, 1991a.

Y. Paulignan, C. MacKenzie, R. Marteniuk, and M. Jeannerod. Selective perturbation of visual input during prehension movements. 1. the effects of changing object position. *Experimental Brain Research*, 83(3):502–512, 1991b.

Y. Paulignan, V. G. Frak, Y. Yoni, and M. Jeannerod. Influence of object position and size on human prehension movements. *Experimental Brain Research*, 114(2):226–234, April 1997.

R. Pelossof, A. Miller, P. Allen, and T. Jebara. An svm learning approach to robotic grasping. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 3212–3218, New Orleans, LA, April 2004. IEEE.

F. Pfeiffer. Grasping with hydraulic fingers-an example of mechatronics. *IEEE/ASME Transactions on Mechatronics*, 1(2):158–167, June 1996.

R. Platt, A. H. Fagg, and R. A. Grupen. Extending fingertip grasping to whole body grasping. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, volume 2, pages 2677– 2682. IEEE, September 2003.

R. Platt, A. H. Fagg, and R. A. Grupen. Reusing schematic grasping policies. In *Proc. Humanoids*, Tsukuba, Japan, 2005.

R. Platt, R. A. Grupen, and A. H. Fagg. Improving grasp skills using schema structured learning. In *Proc ICDL*, Bloomington, 2006.

A. P. del Pobil. Research benchmarks v2. Technical report, EURON report, 2006.

N. S. Pollard. Closure and quality equivalence for efficient synthesis of grasps from examples. *International Journal of Robotics Research*, 23(2):595–614, June 2004.

N. S. Pollard. Synthesizing grasps from generalized prototypes. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, volume 3, pages 2124–2130, Minneapolis, Minnesota, April 1996. IEEE.

J. Ponce, S. Sullivan, J.-D. Boissonnat, and J.-P. Merlet. On characterizing and computing three- and four-finger force-closure grasps of polyhedral objects. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 821–827, Atlanta, Georgia, May 1993. IEEE.

I. Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog, Stuttgart, 1973.

H. Ritter. Neo/nst the graphical simulation toolkit. http://www.techfak.uni-bielefeld.de/ags/ni/projects/neo/neo_e.html.

H. Ritter, J. Steil, C. Nölker, F. Röthling, and P. McGuire. Neural architectures for robotic intelligence. *Reviews in the Neurosciences*, 14(1-2):121–143, 2003.

A. Roby-Brami, N. Bennis, M. Mokhtaria, and P. Baraduc. Hand orientation for grasping depends on the direction of the reaching movement. *Brain Research*, 869(1-2):121–129, June 2000.

A. C. Roy, Y. Paulignan, M. Meunier, and D. Boussaoud. Prehension movements in the macaque monkey: effects of object size and location. *Journal of Neurophysiology*, 88(3):1491–1499, September 2002.

M. Saling, J. Alberts, G. E. Stelmach, and J. R. Bloedel. Reach-to-grasp movements during obstacle avoidance. *Experimental Brain Research*, 118(2):251 – 258, January 1998.

J. K. Salisbury. *Recent advances in robotics*, chapter Kinematic and force analysis of articulated hands, pages 131 – 174. John Wiley and Sons, Inc., New York, NY, USA, 1985.

M. Santello and J. F. Soechting. Matching object size by controlling finger span and hand shape. *Somatosensory and Motor Research*, 14(3):203–212, July 1997.

G. Schlesinger. Der mechanische aufbau der künstlichen glieder. In M. Borchardt et al, editor, *Ersatzglieder und Arbeitshilfen für Kriegsbeschädigte und Unfallverletzte*, pages 321–699. Springer-Verlag, Berlin, 1919.

C. M. Schneck and A. Henderson. Descriptive analysis of the developmental progression of grip position for pencil and crayon control in nondysfunctional children. *American Journal of Occupational Therapy*, 44(10):893–900, October 1990.

J. Schröder, K. Kawamura, T. Gockel, and R. Dillmann. Improved control of a humanoid arm driven by pneumatic actuators. In *Int. Conf. on Humanoid Robots*, Karlsruhe, 2003.

S. Schulz, C. Pylatiuk, and G. Bretthauer. A new ultralight anthropomorphic hand. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 2437–2441, Seoul, Korea, 2001. IEEE.

S. Schulz, C. Pylatiuk, A. Kargov, R. Oberle, and G. Bretthauer. Progress in the development of anthropomorphic fluidic hands and their applications. In *MechRob*, pages 936–941, Aachen, September 2004. IEEE.

The Shadow Robot Company. The shadow dexterous hand, 2006. URL `http://www.shadowrobot.com/hand/overview.shtml`.

M. Simoneau, J. Paillard, C. Bard, N. Teasdale, O. Martin, M. Fleury, and Y. Lamarre. Role of the feedforward command and reafferent information in the coordination of a passing prehension task. *Experimental Brain Research*, 128(1-2):236–242, September 1999.

J. Smeets and E. Brenner. Independent movements of the digits in grasping. *Experimental Brain Research*, 139(1):92–100, 2001.

C. Sollerman and A. Ejeskar. Sollerman hand function test. a standardised method and its use in tetraplegic patients. *Scandinavian Journal of Plastic Reconstructive Surgery*, 29(2):167–176, June 1995.

W. M. Spears. Crossover or mutation. In *Proceedings of the Foundations of Genetic Algorithms Workshop*, pages 221–237. Morgan Kaufmann, 1992.

S. A. Stansfield. Robotic grasping of unknown objects: A knowledge-based approach. *Int. Journal of Robotics Research*, 10(4):314–326, August 1991.

J. F. Steffen. Erfahrungsbasierte Greifstrategie fuer Antropomorphe Roboterhaende. Master's thesis, Faculty of Technology, Bielefeld University, 2005.

J. Steil, F. Röthling, R. Haschke, and H. Ritter. Learning issues in a multi-modal robot-instruction scenario. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, volume Workshop on "Robot Programming Through Demonstration", Oct 2003.

J. Steil, F. Röthling, R. Haschke, and H. Ritter. Situated robot learning for multi-modal instruction and imitation of grasping. *Robotics and Autonomous Systems*, Special Issue on "Robot Learning by Demonstration"(47):129–141, 2004.

I. Steuwer. Aufgabenorientierte qualitätsbewertung mehrfingriger griffe in physikbasierter simulation. Master's thesis, Faculty of Technology, Bielefeld University, 2003.

T. Supuk, T. Kodek, and T. Bajd. Estimation of hand preshaping during human grasping. *Medical engineering & physics*, 27(9):790–797, November 2005.

C. L. Taylor and M. D. Schwarz. The anatomy and mechanics of the human hand. *Artificial limbs*, 2:22–35, May 1955.

P. A. Tipler and G. Mosca. *Physics for scientists and engineers*. Elsevier, Spektrum Akad. Verl., Munich, 2., germ. edition, 2004.

G. Tonietti and A. Bicchi. Adaptive simultaneous position and stiffness control for a soft robot arm. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 1992–1997, 2002.

W. T. Townsend. Mcb - industrial robot feature article - barrett hand grasper. *Industrial Robot: An International Journal*, 27(3):181–188, 2000.

B. Vanderborght, B. Verrelst, R. Van Ham, J. Vermeulen, J. Naudet, and D. Lefeber. Control architecture of lucy, a biped with pneumatic artificial muscles. In *Proc. CLAWAR*, Madrid, 2004.

P. L. Weir, C. L. MacKenzie, R. G. Marteniuk, S. L. Cargoe, and M. B. Frazer. The effects of object weight on the kinematics of prehension. *Journal of Motor Behavior*, 23(3):192–204, September 1991.

A. M. Wing, A. Turton, and C. Fraser. Grasp size and accuracy of approach in reaching. *Journal of Motor Behavior*, 18(3):245–260, September 1986.

Y. J. Wong and I. Q. Whishaw. Precision grasps of children and young and old adults: individual differences in digit contact strategy, purchase pattern, and digit posture. *Behavioural Brain Research*, 154(1):113–123, September 2004.

D. Wren and R. Fisher. Dextrous hand grasping strategies using preshapes and digit trajectories. In *Proc. IEEE Int. Conf. on Systems, Man and Cybernetics*, pages 910–915, Vancouver, BC, Canada, August 1995.

J. E. Yakimishyn and J. Magill-Evans. Comparisons among tools, surface orientation, and pencil grasp for children 23 months of age. *American Journal of Occupational Therapy*, 56(5):564–572, September 2002.

Y. Yokokohji, N. Muramori, Y. Sato, T. Kikura, and T. Yoshikawa. Design and path planning of an encountered-type haptic display for multiple fingertip contacts based on the observation of human grasping behavior. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, pages 1986–1991, New Orleans, LA, USA, April 2004. IEEE.

F. T. J. M. Zaal and R. J. Bootsma. Accuracy demands in natural prehension. *Human Movement Science*, 12(3):339–345, May 1993.

Glossary

approach distance d Distance between the pre-grasp position and the grasp position,

both strategy A contact strategy in which the first contacts with the object are made by the thumb and any opposing finger simultaneously,

contact simultaneity The grasping fingers touch an object at the same point in time,

contact strategy Determines the chronological sequence of opposing finger contacts with the object during grasp closure. Three different contact strategies are defined: both strategy, finger drag strategy, and thumb push strategy,

default grasp Grasp before applying the optimisation strategy (similar: *default pre-grasp* and *default target grasp*),

DOF degree (or degrees) of freedom,

finger drag strategy A contact strategy in which any finger except the thumb touches the object first,

GOT% Grasp forming time in percentage of grasp closure time,

GOT% Grasp opposing time in percentage of grasp closure time,

grasp g A grasp g identifies a pre-grasp, a target grasp, an orientation $\boldsymbol{o}$, a position $\boldsymbol{p}$, and an approach distance d,

grasp closure time (GCT) Time between the points of peak hand opening and when the last finger touches target object,

grasp forming time (GFT) Time between the points when the first and the last grasping finger touches the target object,

grasp opposing time (GOT) Time between the points when the first finger and the first opposing finger touches the target object,

grasp position Position of the hand for grasp closure, determined by position $\boldsymbol{p}$ and associated orientation $\boldsymbol{o}$,

grasp posture Pre-grasp posture or target grasp posture,

grasp strategy Sequence of steps to apply a grasp as defined in Section 2.2.4,

grasp type t The grasp type t defines the number of grasping fingers and qualitatively determines the target touch areas of the hand. The requirements on grasp types t_1 to t_4 are described in Section 2.3.2. To *realise* a grasp type means: to define a grasp in a robot hand setup that fulfils the requirements,

grasping / to grasp / to apply a grasp Execution of the grasp strategy,

grasping finger Finger that is expected to touch the target object when the grasp is applied,

grip posture Posture a hand adopts at the end of the grasp closure when the fingers touch the target object,

hand frame H As defined in Section 3.2.2,

new grasp Grasp after applying the optimisation strategy (similar: *new pre-grasp* and *new target grasp*),

object zero-frame O_z As defined in Section 3.2.2,

orientation $\boldsymbol{o}$ Rotation vector $\boldsymbol{o} = (\gamma, \beta, \alpha)^{-1}$ between object zero-frame O_z and robot hand frame H determining the orientation of the hand in pre-grasp position and in gras position,

position $\boldsymbol{p}$ Translation vector $\boldsymbol{p} = (x, y, z)^{-1}$ between object zero-frame O_z and hand frame H determining the grasp position,

pre-grasp Set of joint angles to pre-shape the hand before the target grasp is applied. To *apply* a pre-grasp means that the joint angles are actuated by the robot hand controller,

pre-grasp position Position to be approached by the hand, before it is moved to the grasp position when being in pre-grasp posture,

pre-grasp posture Posture a hand adopts when applying the pre-grasp,

standard grasp A *standard grasp* is a grasp g that fulfils the requirements of a grasp type t and is realised on a robot hand setup in a developmental process as described in Section 3.2.3,

target grasp Set of hand joint angles for taking hold of an object. To *apply* a target grasp means that the joint angles are actuated by the robot hand controller,

target grasp posture Posture a hand adopts when executing the target grasp without being obstructed by an object,

thumb push strategy A contact strategy in which the thumb touches the object before any other finger,

time frame t_f Contacts are considered to be simultaneous if they occur in a time frame specified by t_f,