

Tradeoffs in Neuroevolutionary Learning-Based Real-Time Robotic Task Design in the Imprecise Computation Framework

Pei-Chi Huang¹, Luis Sentis², Joel Lehman³, Chien-Liang Fok², Aloysius K. Mok¹, and Risto Miikkulainen¹

¹ Department of Computer Science, University of Texas at Austin

² Department of Mechanical Engineering, University of Texas at Austin

³ IT University of Copenhagen

Technical Report

November, 2017

Abstract

A cyberphysical avatar is a semi-autonomous robot that adjusts to an unstructured environment and performs physical tasks subject to critical timing constraints while under human supervision. The paper first realizes a cyberphysical avatar that integrates three key technologies: body-compliant control, neuroevolution, and real-time constraints. Body-compliant control is essential for operator safety because avatars perform cooperative tasks in close proximity to humans; neuroevolution (NEAT) enables “programming” avatars such that they can be used by non-experts for a large array of tasks, some unforeseen, in an unstructured environment; and real-time constraints are indispensable to provide predictable, bounded-time response in human-avatar interaction. Then, we present a study on the tradeoffs between three design parameters for robotic task systems which must incorporate at least three dimensions: (1) the amount of training effort for robot to perform the task, (2) the time available to complete the task when the command is given, and (3) the quality of the result of the performed task. A tradeoff study in this design space by using the imprecise computation as a framework is to perform a common robotic task, specifically, grasping of unknown objects. The results were validated with a real robot and contribute to the development of a systematic approach for designing robotic task systems that must function in environments like flexible manufacturing systems of the future.

1 INTRODUCTION

Although modern robots can perform complex tasks competently through hand-designed algorithms [48], it remains challenging [43] to create robots capable of completing mission-critical tasks in unstructured environments without complete dependence upon a human (e.g. through teleoperation). Our framework to tackle this challenge is based upon the concept of a “cyberphysical avatar,” defined to be a semi-autonomous remote robotic system that adjusts to an unstructured environment and performs physical tasks subject to real-time constraints under human supervision [7].

The key scientific question of interest is what form and minimum degree of human supervision is required to enable a robot to perform a particular type of task. To answer this question, it is important to recognize the fact that the quality of robotic task performance is a function of at least two parameters: the amount of training the robot has had through machine learning algorithms, as well as the tightness of the real-time task deadline that the robot is to meet. For example, if we give the robot one second to grasp an unknown object, it is likely that the grasp will not be as firm and reliable as what the robot would be able to achieve given ten seconds to complete the grasp. The goal of our research is to perform a systematic investigation of the tradeoffs between the training effort, the resulting quality of the robotic task, and the time the robot takes to perform the task. Understanding this tradeoff is essential to design robots that can function effectively in real time.

The specific robotic task we use for this investigation is the *grasping* of an unknown object by the robot. While robotic grasping has received significant research attention, the type of tradeoff investigation mentioned above is not. We use *Dreamer*, a humanoid torque-controlled mobile robot as our experimental platform. For training *Dreamer* to perform the grasping task, we adopt the NEAT (NeuroEvolution of Augmenting Topologies) machine learning method. The design space of the robotic grasping task has three dimensions: (1) The training effort, measured by the time used in running the NEAT algorithm to train *Dreamer* to perform the grasp. (2) The task completion time, defined by the time *Dreamer* has to perform the grasping task. (To enforce the completion time constraint, a trajectory planner is used to compute the way-points for the trajectory that connects the initial and final configuration of *Dreamer* within the target completion time. The actual physical trajectory is realized by *Dreamer*’s on-board controller). (3) The quality of the grasp is evaluated both in simulation and on a physical robot. Our training method relies on simulations modeled with GraspIt! [31], which is an open-source grasp simulation environment that models the targeted robotic hand and can evaluate grasp quality. In this paper, we focus on the tradeoff between training effort and amount of time allowed to attempt the grasp. The metric for evaluating this tradeoff is the quality of the resulting grasp. Implicit in the premise of this paper is that training time is a limited resource. In traditional process automation areas such as manufacturing, the design of the system may often take months, or even years, because the system is then expected to run continuously without further significant modification. However, many future robotic applications require much faster training. For example, for flexible manufacturing a robot may be trained to assemble only a small number of a particular product (i.e. rapid small-lot manufacturing) before being retrained to make a different product. It is therefore important to understand how much training is enough for the robot to satisfactorily perform a task. Our investigation is also relevant in light of recent successes in deep learning [20, 49]. Similar to the neuroevolution approach applied here, successful application of deep learning involves multi-layer neural networks, and depends upon a well-chosen network topology and a sufficient number of training examples. Thus, an important question is how task performance improves as a function of the number of training cases. A plausible conjecture is that with a properly connected multi-layer network, the performance curve may exhibit fast

improvement once past some critical number of input training cases. For example, a juggling robot may ‘suddenly’ acquire the juggling skill once some basic hand-eye coordination ‘invariant’ has been captured by the evolving neural network. To answer this type of question, the tradeoff study in this paper should be useful. In particular, we use a framework from the area of real-time systems research called the imprecise computation model [25, 26], to help explore the boundary region of tolerance and find best effort techniques. Our work can be viewed as providing a realistic basis for some of the scheduling work done by the real-time systems community in the past two decades.

The remainder of this paper is organized as follows. Section 2 reviews related work and describes the neuroevolution method and the imprecise computation. Section 3 describes system integration and its architecture. Section 4 introduces our model of the whole-body compliant grasp of *Dreamer* and its hierarchical control structure. The learning approach is presented in Section 5, while Section 6 describes experimental results and their evaluation. Section 7 measures the grasping performance by applying imprecise computation. Finally, Section 8 and Section 9 conclude by reviewing remaining problems and future work.

2 BACKGROUND AND RELATED WORK

This section reviews previous machine learning approaches to robotic grasping in 2.1, the neuroevolution method applied in the experiments in 2.2, the imprecise computation technique to analyze performance in 2.3, and grasp quality measurement in 2.4.

2.1 Robotic Grasping through Machine Learning

Impressive progress has been made in learning to grasp novel objects [19, 32, 37, 38, 42–44]. To the best of our knowledge, previous methods use only simple hand models [32, 43] and are not directly applicable for the target hand in this work (i.e. *Dreamer* robot’s hand, *Mekahand*). Also, transferring controllers from simulation to reality remains challenging [14, 23].

Related to the approach described here are previous artificial neural networks (ANNs) approaches that simulate arm kinematics [34, 36, 39]. Other approaches use reinforcement learning techniques [17, 55] to explore search spaces optimally for control strategies, and learning from demonstration (LfD) [1] to improve grasping capability, and partially observable Markov decision processes (POMDP) [9, 10] to choose optimal control policies. Yet only a few of these methods were tested in the real world.

2.2 Neuroevolution

Neuroevolution (NE) is an approach where an evolutionary algorithm is applied to learn the structure of an ANN, its connection weights, or both [51]. Compared with other machine learning methods, neuroevolution is unique in two main ways.

First, most other learning methods are supervised, i.e., they learn behavior that approximates a given set of examples [8]. It is important that such examples are carefully chosen to ensure that the training process results in learning a function that smoothly interpolates between them. For instance, in robotic grasping, a training set consists of grasping situations paired with the corresponding optimal grasping behavior. Because optimal behavior is often not known, it is unclear how such examples can be produced to cover representative situations well. In contrast, neuroevolution is a reinforcement learning method, and as such it does not require training examples where ideal

behavior is known. Second, neuroevolution does not rely on complete state information. Other methods that are designed to learn under sparse reinforcement, such as Q-learning (or value function learning in general) often assume that the current state of the system is completely known [53]. However, if objects are occluded or situation varies dynamically, it is difficult for such methods to differentiate between possible situations because the observed values of actions cannot be associated with the correct underlying state. Neuroevolution solves the problem by evolving recurrent connectivity; recurrence establishes memory that make it possible to distinguish between states.

One complication in applying neuroevolution to a complex domain like robotic grasping is that the ideal network topology (i.e. how many neurons compose the network and how are they interconnected) is not a known *a priori*. Because the depth image input contains many low-level features (i.e., pixels), a fully connected network with many hidden neurons may have an intractable number of parameters to tune. This motivates the NeuroEvolution of Augmenting Topologies (NEAT; [51]) method which is a popular method for evolving both network topology and connection weights. With NEAT, the ideal network topology needs not be known *a priori*, but is discovered automatically as part of evolution. The NEAT has shown previous promise both in grasping [2] and in extracting features from low-level input [18].

2.3 Imprecise Computation

In real-time applications (e. g., safety-critical applications), it is difficult for every critical task to meet its deadline. Imprecise computation is a scheduling technique that reduces the amount of time used on a job by means of sacrificing levels of quality of service (QoS) [25, 26]. If the best desired quality of results cannot be obtained, imprecise computation decreases the QoS to make it possible to meet timing constraints of real-time tasks while still keeping the quality within an acceptable range. When the system can not produce accurate results in a timely manner, the graceful degradation can be achieved by providing users with an approximate quality of acceptable results to prevent timing failures. For the real-time robotic task design in the imprecise computation framework, consider a grasping task, implementing two parts. A mandatory part that the task must complete before its deadline to achieve the minimum expected quality, and an optional part that can be discarded at any time to improve the quality results by associating a reward with their execution.

2.4 Grasp Quality Measurement

Given an object, finding a suitable grasp configuration among the infinite set of candidates has been studied extensively in the robotics community over the two decades [5, 32, 41, 47, 50]. The algorithms of producing a feasible grasp require determining proper dexterous hand configurations (i.e., *Mekahand*) as well as contact points on the objects. Much previous grasping quality research focuses only on contact types and positions, ignoring hand geometry and kinematics. Other measures assume simple grippers. Roa and Suárez [41] reviewed and analyzed the performances of evaluation methodologies of grasp quality in the literature. One of the most popular quality measure approaches [41] is to consider the force constraints, a grasping quality measure is defined as the externally largest perturbation wrench where the grasp must be capable of resisting in any direction of forces ([6, 16]); the frictional grasps of mathematical basis has been proposed [30, 33, 54], and is also applied in the work on predicting grasping [13, 31, 35]. A grasp is in balance when the summation of forces and torques is null and the fingers is applied in the object (external disturbances) [3, 15, 24, 27]. The method of grasping quality is based on an analytic formulation to compute how friction occurring grasping acting on those contacts affects the space of forces and

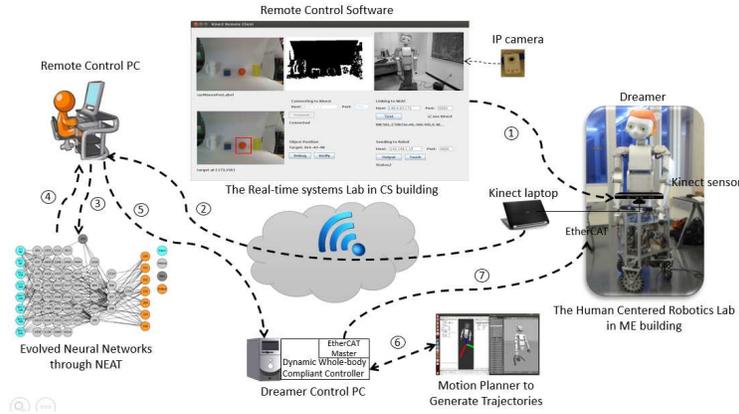


Figure 1: The semi-autonomous robotic system. (1) A human supervisor connects to the Kinect laptop, (2) captures a depth image, and (3) parses the depth array to serve as input to an evolved Artificial Neural Network (ANN). (4) The neural network’s output is interpreted as directions to control *Mekahand*’s position and orientation, and is sent to the supervisor. (5) The supervisor sends commands to manipulate *Dreamer* robot. (6) Motion planner generates a trajectory from the initial state to the final state. (7) The controller of the wheeled humanoid avatar controls its body and arm to destination in unstructured environments. The conclusion is that the system integrates real-time vision, neuroevolution as a training method, and control manipulator while skillfully reaching an object through the man-machine interface.

torques that can be applied to an object and further rank/decide what set of contacts are appropriate to grasp. This method was adopted in GraspIt! to measure grasp quality of the *Mekahand*, which is used in this paper.

The architecture of the system is described in the next section.

3 THE SEMI-AUTONOMOUS ROBOTIC SYSTEM

Having summarized the motivation for designing cyberphysical avatars, emphasizing the important contributions that they could make, we turn to present the actual architecture of a cyberphysical avatar, also called a semi-autonomous robotics system, used interchangeably in this paper.

This semi-autonomous robotic system comprises a mobile dexterous humanoid robot *Dreamer* with its whole body control system, and devised machine learning algorithms (Neuroevolution) including awareness of the environment complexity and sensing unpredictable world, and a real-time physical distribution network, and a series of cost-effective, real-time and vision system. The specific task explored in this work is controlling the *Dreamer* robot to approach and pick up a designated target object under remote human supervision in a real-time environment. The physical realization of the cyberphysical avatar has been implemented in the Human Centered Robotics Laboratory (HCRL) at the University of Texas at Austin, and the portable remote control user interface is located in another building nearby.

Figure 1 illustrates an overview of the semi-autonomous robotic system. *Dreamer* consists of a torso, two arms, two hands, an anthropomorphic head [46]. The *Dreamer* is equipped with torque and sensors to provide force compliant capabilities. A desktop PC running Ubuntu Linux with the RTAI Real-time Kernel executes the models and control infrastructure to govern *Dreamer*’s behavior via EtherCAT serial ports. Two types of cameras are installed in the system. A Kinect camera connects to a laptop and is installed in front of the robot to capture images and depth information,

and an IP camera is installed at the ceiling to capture *Dreamer's* surrounding environment. The Kinect laptop connects to the avatar and sends images to the remote supervisor.

A grasping experiment is achieved as follows. First, the human supervisor directs the *Dreamer* robot with a command to grasp the desired object. The cyberphysical avatar communication software relays the human input and depth information to a neural network that has been evolved with NEAT. Recall that NEAT's role is to train a neural network in a simulator to produce the appropriate outputs for *Dreamer* to act on. To apply NEAT to learn where and how to grasp an object requires both training scenarios and a measure for evaluating performance. GraspIt! [31] provides the interactive simulation, planning, analysis, and visualization. The neural network (trained off-line) outputs the appropriate positions and orientations to *Dreamer* robot which then moves towards the destination and grasps the targeted object with its *Mekahand*.

Next, we introduce the designed dynamic controller of *Dreamer* and the *Mekahand* model.

4 DYNAMIC CONTROL OF HUMANOID ROBOTS IN UNSTRUCTURED ENVIRONMENTS

Dreamer's upper body consists of 3-dof torso, 7-dof arms and a 12-dof *Mekahand*, as shown in Figure 2. The 3-dof torso has one unactuated joint which is coupled with the waist joint. The hand also has five actuated joints and seven coupled unactuated joints, shown in Figure 2. To simplify the controller, we divided the controller into one for controlling the body and the arm, and the other for controlling the hand.

To control the body and the arm together, skill modeling and dynamic control of the robot are necessary. The prioritized whole-body compliant controller (WBC) is used for our purpose [46]. In WBC, first an objective is set and then a task is defined by a Jacobian [45] to derive the relations between the robot's 10-dimensional joint spaces and the M-dimensional operational space. The controller is derived from the following constrained system dynamics equations.

$$A\ddot{q} + b(q, \dot{q}) + g(q) + J_c^T \lambda = U^T T, \quad (1)$$

where A is the mass matrix of the system, q is the joint coordinate vector, b is the torque caused by Coriolis and Centrifugal effects, g is the torque caused by gravity, J_c is the constrained Jacobian, λ is the Lagrangian multiplier that describes the constrained joints, U is the actuation matrix, and T is the torque input to the system. The reason why the constrained Jacobian and the Lagrangian multipliers are shown in the system is to model the underactuated torso and the transmission constraint. The body joints 1 and 2 are coupled together. Therefore, we can specify the constraint as follows:

$$\dot{q}_1 - \dot{q}_2 = 0, \quad (2)$$

$$J_c \dot{q} = 0, \quad (3)$$

$$J_c = \begin{bmatrix} 0 & 1 & -1 & 0 & \dots & 0 \end{bmatrix} \in \mathbf{R}^{1 \times 10}. \quad (4)$$

We can take the constrained mass matrix Λ_c , the dynamically consistent generalized inverse of J_c ,

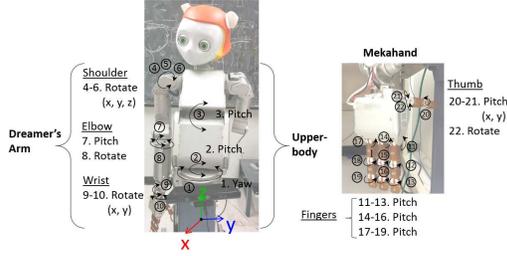


Figure 2: *Dreamer's* upper body and the *Mekahand*. *Dreamer* contains 3-dof torso (1-3), a 7-dof arm (4-10) and a 12-dof *Mekahand* (11-22). Each unactuated/actuated joint is coupled with another joint. The conclusion is that since many DOFs increase in difficulty of *Dreamer's* balance control while grasping, it is necessary to design a skill modeling and dynamic control of *Dreamer*.

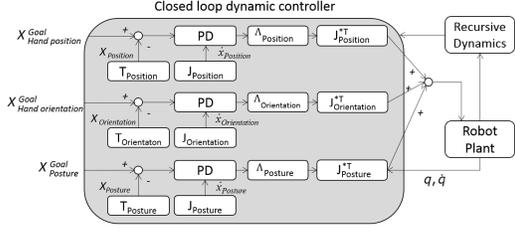


Figure 3: The designed whole-body compliant controller (WBC). The tasks of hand position, hand orientation and the posture of *Dreamer* upper-body are combined to perform a grasping skill. The feedback/feedforward control policies contribute to the closed-loop dynamic controller. The conclusion is that the designed control structure can effectively utilize dynamic and contact models of the physical robot in unstructured environment.

and the constrained null space N_c to derive the constrained dynamic equation as follows:

$$\Lambda_c \triangleq \left(J_c A^{-1} J_c^T \right)^+ , \quad (5)$$

$$\bar{J}_c \triangleq A^{-1} J_c^T \Lambda_c , \quad (6)$$

$$N_c \triangleq I - \bar{J}_c J_c , \quad (7)$$

$$\ddot{q} = A^{-1} N_c^T U^T T . \quad (8)$$

Then, we can define task space specifications to derive the desired forces in the constrained dynamic systems. In the case of the position task that makes the end-effector (hand) approach the object, the task Jacobian is defined as

$$\dot{x} = J_{\text{position}} \dot{q} , \quad (9)$$

where x is the end-effector coordinate. The task Jacobian can describe the relation between the joint velocities and the coordinate system that a supervisor expects to control. The task Jacobian does not include the constrained dynamics, so we need to project this Jacobian to the constrained space and then generate the constrained task Jacobian,

$$J_{\text{position}}^* \triangleq J_{\text{position}} \overline{UN_c} . \quad (10)$$

The grasping skill, including posture, position and orientation, is defined as a juxtaposition of multiple operational tasks to help translate between high-level goals, such as those provided by the planning algorithms, and the operational tasks. In the robot's environment, a skill is composed of the three tasks in Figure 3: hand position, hand orientation, and the posture of *Dreamer's* whole body. The control structure can be expressed as

$$\tau_{\text{control}} = J_{\text{position}}^{*T} F_{\text{position}} + J_{\text{ori}}^{*T} F_{\text{ori}} + J_{\text{posture}}^{*T} F_{\text{posture}} , \quad (11)$$

where F_{position} , F_{ori} and F_{posture} are the force or impedance commands to control the hand, and J_{position}^{*T} , J_{ori}^{*T} and J_{posture}^{*T} are the whole-body task Jacobians [45]. The grasping process consists of

three behaviors, among which the *Mekahand* position and orientation have higher priority than the whole-body posture task because the latter is of secondary concern in the context of grasping.

The feedback control policies for the entire controller are shown in Figure 3, which depicts the closed-loop dynamic controller. In the figure, the Goal Hand Position/Orientation is a compliant hand position/orientation that enables the robot to reach a designated position/orientation. The Goal Posture control exploits the remaining DOFs to stabilize self-motions. The proposed feedback/feedforward control laws are

$$F_{\text{position}} = \Lambda_{\text{position}}^* (-k_{p, \text{position}} e_{\text{position}}^{\text{goal}} - k_{v, \text{position}} \dot{x}_{\text{position}}) + p_{\text{position}}, \quad (12)$$

$$F_{\text{orientation}} = \Lambda_{\text{ori}}^* (-k_{p, \text{ori}} e_{\text{ori}}^{\text{goal}} - k_{v, \text{ori}} \dot{x}_{\text{ori}}) + p_{\text{ori}}, \quad (13)$$

$$F_{\text{posture}} = \Lambda_{\text{posture}}^* (-k_{p, \text{posture}} e_{\text{posture}}^{\text{goal}} - k_{v, \text{posture}} \dot{x}_{\text{posture}}) + p_{\text{posture}}, \quad (14)$$

where $\Lambda_{\text{position}}^*$, Λ_{ori}^* and $\Lambda_{\text{posture}}^*$ are the inertial matrices projected in the manifold of the constraints, $e_{\text{position}}^{\text{goal}}$, $e_{\text{ori}}^{\text{goal}}$ and $e_{\text{posture}}^{\text{goal}}$ are feedback error functions, k_p , k_v are gain matrices, and p_{position} , p_{ori} and p_{posture} are gravitational terms. This structure is a derivation of the previous work on compliant whole-body control [46].

Since our designed control structure can effectively use dynamic and contact models of the physical robot in its environments, it is able to optimize the process of approaching and grasping objects simultaneously, and to achieve precise tracking of forces and trajectories within the contact conditions. Thus, the grasping skill is acquired through neural network described next.

5 ACQUIRING GRASPING SKILLS THROUGH NEUROEVOLUTION

Although robots can be often controlled through carefully hand-designed algorithms, this paper explores one way in which reinforcement learning methods can provide a significant advantage: optimization of robot behaviors. The difficulty in designing effective control algorithms by hand suggests that machine learning may be a desirable approach, yet to apply common supervised learning algorithms requires a corpus of labeled examples. However, in contrast to supervised learning, a measure of quality is sufficient to apply reinforcement learning algorithms. In many cases, it is much easier to derive a measure of how desirable a particular behavior is than it is to either hand-construct that behavior or provide a comprehensive corpus of optimal example behaviors.

For this reason, the approach in the paper applies reinforcement learning to facilitate learning high-level behaviors that can be then invoked by a human operator. In particular, neuroevolution algorithms have proven effective in domains with low-level continuous features that are characteristics of the problem here, i.e. learning to grip objects given depth sensor information. Thus, this section introduces our approach [11], which is based on applying a popular neuroevolution method called NEAT to the GraspIt! simulation environment. Subsection 5.1 introduces the grasping learning approach; Subsection 5.2 then describes the learning process, specifying the input and output layers, as well as fitness function, and also how to speedup the process. Finally, Subsection 5.3 presents the tradeoffs evaluation in the imprecise computation framework.

5.1 Grasp Learning Approach

Our approach takes inspiration from Kohl et al. [18] who showed that neuroevolution can develop effective automobile warning systems from only low-level sensor input (i.e. pixels) taken from a digital camera. A similar vision-based feature extraction approach is applied here, where through neuroevolution the *Mekahand* robotic arm learns appropriate hand positions and orientations for grasping. Such learning is enabled by interacting with objects in the GraspIt! simulation environment, which is described next followed by the approach to measure grasping quality and determine a visual bounding box for grasping.

5.1.1 GraspIt! Simulation Implementation

To apply neuroevolution to learn where and how to grasp an object requires both training scenarios and a metric for evaluating performance. GraspIt! [4,29] facilitates simulating the *Mekahand* robot in representative grasping tasks and aids in measuring the quality of resulting grips.

GraspIt! only provides a rough *Mekahand* model, so we extended the simulator to better model it. In GraspIt!, the *Mekahand* is defined by one DOF for each knuckle in each finger, with an additional DOF for the thumb’s rotator. The mechanics of this model are modified here to augment two aspects of the simulation. First, controlling the wrist is not modeled by default, but is an important DOF. Therefore, a wrist component was added to the *Mekahand* model supplied by GraspIt!. Second, most of the DOFs in the real *Mekahand* are not actuated, although they are modeled as actuated in the GraspIt! simulation. Each finger of the real *Mekahand* consists of three joints that are all connected by a single rubber tendon. Thus when the finger curls, all three knuckles curl in unison. Therefore, the torques in GraspIt! were adjusted such that the set of torques given to a single finger are equivalent to the torques initiated by stretching the rubber tendon in the real robot.

GraspIt! uses a quaternion to represent the rotation of a 3D object. Since our learning output applies axis-angle representation in a 3D Euclidean space. Our implementation automatically translates the quaternion into the axis-angle representation in a 3D Euclidean space for the output.

5.1.2 Measuring Grasp Quality

An evolutionary search optimizes a fitness function that measures the quality of candidate solutions. Because robust grasping behaviors are desired in this experiment, an important consideration is how to measure the quality of grasps appropriately. Recall that the approach for grasp measure [3,15,24,27,30] was applied in GraspIt! to measure grasp quality of the *Mekahand*. Given a 3D object and posture of the *Mekahand*, their measure can accurately identify the types of contact points between the links of the hand and the object and compute the grasp’s quality. Figure 4 illustrates the score for each grasp when applied to different objects (a single cylinder, sphere, cube, and mug) from the different positions and orientations in ascending order. The grasp quality is -1 if the *Mekahand* just touches the object; otherwise, the quality is larger than zero. This grasp quality metric can yield different scores, which can be used to score each grasp for machine learning.

The next Subsection shows how this approach can be applied to a human-supplied bounding box and focusing the robot’s visual processing on the target object, thereby lessening the dimensionality of the robot’s computer vision processing.

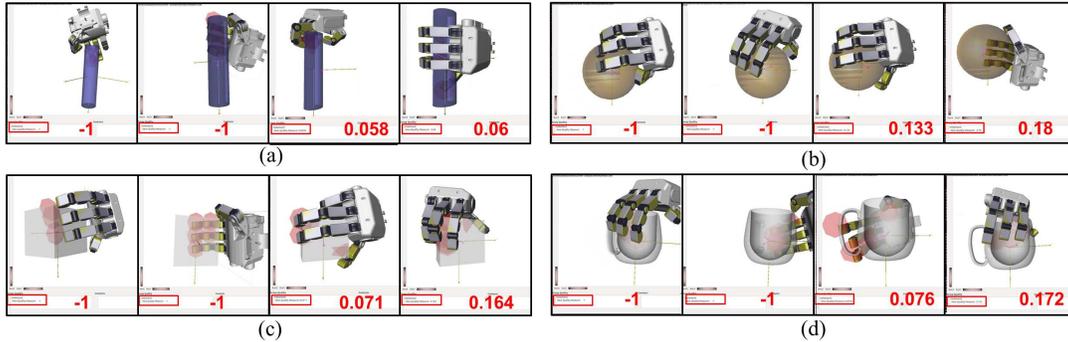


Figure 4: Measuring grasp quality. One grasp acting on a single (a) cylinder, (b) sphere, (c) cube, (d) and mug is represented by different scores based on the *Mekahand*'s position and orientation. The grasp quality is -1 if the *Mekahand* just touches the object; otherwise, the quality is larger than zero. This quality metric can be utilized to score each grasp for machine learning.

5.1.3 Visual Bounding Box

In the experiment, ANNs through exploration learn how to grasp objects by integrating information from a high-dimensional depth image provided by a Kinect sensor. To better focus on the most important features of the depth image, a bounding box strategy was implemented. For each object extracted from the original scene, image data was considered only from within a supervisor-specified bounding box. The bounding box thus serves to minimize the number of irrelevant pixels considered and then simplifies the learning problem.

The training process with the bounding box method proceeds as follows. GraspIt! loads a scene, and then two mouse clicks from the user specify a rectangular bounding box that encompasses the object. In the simulated implementation, because all relative 2D coordinates of each object can be determined, an encompassing bounding box is automatically generated and centered on the desired object. For simplicity, all the computed bounding boxes have the same size. The boundary range can be mapped to four coordinates. For example, in Figure 5, a cube is chosen, so the bounding box is $(C_x, C_y), (C'_x, C_y), (C_x, C'_y), (C'_x, C'_y)$. The depth array of the bounding box is then divided into $M \times N$ pixels that are given to the ANN being evaluated as input data.

To simplify the implementation, the position of the camera sensor is always set such that the origin $O_{3d} (0, 0, 0)$ in the GraspIt! scene is in the center of the 2D plane, as shown in Figure 5. Because the input is reduced to a small part of the overall depth image, after the ANN produces the output, the position of each object must be offset relative to the bounding box. For example, in Figure 5, for the cube, Δx and Δy should be added to the position of the output, for mapping to the normalized origin position.

5.2 Learning Process

In learning process, we elaborate the NEAT method first. Then, combining neuroevolution with the grasping task requires specifying the input and output layers of the neural network, as well as a fitness function to evaluate grasps. A schematic description of the general framework combining GraspIt! and NEAT is depicted in Figure 6. Note that there are no supervised examples that the algorithm attempts to emulate. Instead, the algorithm learns from reinforcement feedback based on only the measured quality of attempted grasps. In this way, evolution can discover solutions that work well even when the optimal behaviors are unknown.

5.2.1 The NEAT Method

Behaviors are evolved for robots that are controlled by artificial neural networks (ANNs). Thus, the NEAT method is suitable to underpin our experiments because it is broadly utilized [21, 22, 40, 51, 52]. NEAT evolves starting from a population of small and simple networks, and *complexifies* the network topology into diverse species over generations, leading to increasingly sophisticated behavior. Here, a brief review of the NEAT methodology was provided; for comprehensive introductions see e.g. [21, 51, 52]. To keep track of which gene is being added to new genes, a *historical marking* is uniquely assigned to each new structural component. During crossover, genes with the same historical markings are aligned, effectively producing meaningful offspring. *Speciation* in NEAT protects structural innovations by decreasing competition among different structures and network complexities, allowing newer and more complex structures to be adjusted. Networks are assigned to species according to the extent to which they share historical markers. *Complexification* resembles how genes are added during natural evolution [28], and is thus supported by speciation and historical markings, thereby giving NEAT to establish high-level features early in evolution. In addition, the ability of NEAT to evolve increasingly complex ANNs is well-suited to robotic grasping behaviors, which need potentially complex evolved structure.

5.2.2 Input and Output Layers of Neural Network

Each ANN evaluated by NEAT receives input data denoting the current state of the robot in its environment. It is thus necessary to encode such state information, which includes the position of the target object as well as information about the object’s shape. To eliminate dependency on high-level human-provided features of the grasped object, the object’s state is described by general low-level features provided by a depth map. In particular, each pixel in the depth information array is assigned a unique input node, as shown in Figure 6. In this way, the network can potentially learn to associate the state of an arbitrary object in an arbitrary environment with an appropriate grasping strategy.

Each ANN predicts where the object is and in what direction to grasp the object by outputting 3D hand positions and orientations. Note that each dimensional coordinate of the *Mekahand*’s position and orientation maps to one output neuron. Because the orientation is expressed in an axis-angle format (e.g. a 3D axis vector and one angle), the total dimensionality is seven, i.e. the ANN has seven output neurons.

Evolution is initialized with ANNs with input nodes that are fully connected to at least a single hidden neuron, and with the hidden node fully connected to the output neurons. Recall that during evolution, ANNs can accumulate additional connections and nodes through structural mutations that augment network topology.

5.2.3 Grasping Fitness Function

A key element of the experimental design is to construct a fitness function to guide the search process for an appropriate ANN grasp controller. The design of a fitness function is a critical factor for guiding successful evolution.

In particular, in this experiment, the fitness of a network n with respect to an object O has four components:

- f_1 : Grasp quality metric Q , described in Subsection 5.1.2.

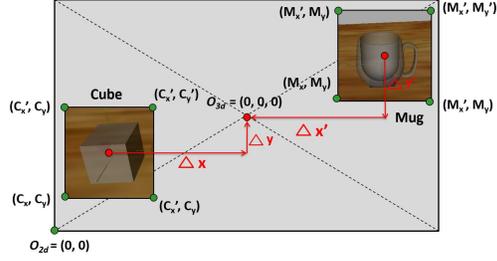


Figure 5: Bounding boxes of a cube and mug, and the output shift offsets Δx and Δy ($\Delta x'$ and $\Delta y'$). Because all relative 2D coordinates of each object are known, an encompassing bounding box is generated centered on the desired object. This figure shows that the boundary range can be mapped to four coordinates. To simplify implementation, the position of the camera sensor is always set such that the origin $O_{3d} (0, 0, 0)$ in the GraspIt! scene is always in the center of 2D plane. The conclusion is that a bounding box strategy can focus on the most important features of the depth image.

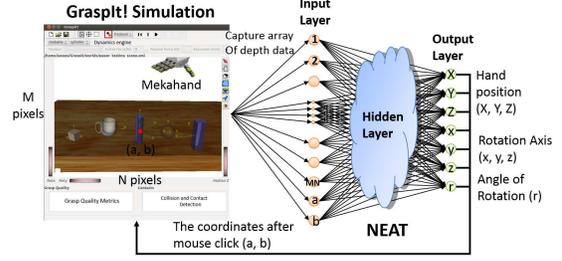


Figure 6: Representation of the designed grasp controller network. The left side of the figure shows GraspIt! simulation environment; the right side of the figure shows a neural network receiving input consisting of depth data and the goal coordinate (a, b) on the GraspIt! visual input scene. The network has seven output nodes: hand position (X, Y, Z) , rotation axis (x, y, z) and rotation angle (r) . Note that NEAT can add internal hidden nodes as evolution progresses. The figure shows how to implement grasping experiments with NEAT in GraspIt!.

- f_2 : The reciprocal of Euclidean distance $d(\vec{P}_i, \vec{O}_i)$ between the hand position computed by the neural network (\vec{P}_i) and a desired object (\vec{O}_i). Note that \vec{P}_i and \vec{O}_i are vectors.
- f_3 : The reciprocal of Euclidean distance $d(\vec{P}_i, \vec{S}_i)$ between the hand position computed by the neural network (\vec{P}_i) and the actual hand coordinate after interacting with the environment (\vec{S}_i). Note that \vec{P}_i and \vec{S}_i are vectors.
- f_4 : An angle θ between the *Mekahand* and grasping object. Let \vec{V}_1 be one vector from the center of the palm to the fingertip of the thumb; let \vec{V}_2 be the vector from the hand position to the center-of-gravity of the desired object; let \vec{V}_3 be the vector indicating the direction of the hand's axis of rotation. Let θ_1 (θ_2 , respectively) be an angle between \vec{V}_1 and \vec{V}_2 (\vec{V}_2 and \vec{V}_3 , respectively). To ensure that the center of palm always turns toward the object, the sum of θ_1 and θ_2 must be roughly around 90° . Figure 7(a) is one good case where the hand axis-angle is almost perpendicular to the object. Figure 7(b) is one bad case where the palm of hand is not orientated toward the object. Here, \vec{V}_1 and \vec{V}_2 (\vec{V}_2 and \vec{V}_3 , respectively) are normalized so that $\|\vec{V}_1\| = \|\vec{V}_2\| = 1$ ($\|\vec{V}_2\| = \|\vec{V}_3\| = 1$, respectively). The angle θ is the sum of θ_1 and θ_2 as follows:

$$\begin{aligned}
 \theta &= \theta_1 + \theta_2 = \theta(\vec{V}_1, \vec{V}_2) + \theta(\vec{V}_2, \vec{V}_3) \\
 &= \frac{\arccos(\vec{V}_1 \cdot \vec{V}_2)}{\|\vec{V}_1\| \|\vec{V}_2\|} + \frac{\arccos(\vec{V}_2 \cdot \vec{V}_3)}{\|\vec{V}_2\| \|\vec{V}_3\|}.
 \end{aligned} \tag{15}$$

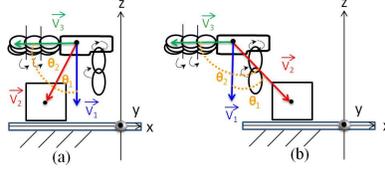


Figure 7: An angle θ between the *Mekahand* and grasping object. \vec{V}_1 is a vector from the center of palm to the fingertip of the thumb; \vec{V}_2 is a vector from the center of palm to the center-of-gravity of the cube; \vec{V}_3 is a vector from the *Mekahand*'s rotation axis. (a) A good case where the palm's center is facing the target object; the sum of θ_1 and θ_2 is almost 90° . (b) A bad case where the palm's center is not facing the target object; the sum of θ_1 and θ_2 is larger than 90° . The conclusion is that because the center of plam facing towards an object can increase the grasping opportunity, the component was added to reward the fitness function.

Thus, the fitness function f of a network n is defined as follows:

$$\begin{aligned} f &= f_1 + f_2 + f_3 + f_4 \\ &= \gamma Q + \frac{\beta}{d(\vec{P}_i, \vec{O}_i) + \alpha} + \frac{\lambda}{d(\vec{P}_i, \vec{S}_i) + \epsilon} + f(\theta). \end{aligned} \quad (16)$$

where $\alpha, \beta, \gamma, \lambda$ and ϵ are constants chosen to balance the various parameters. Note that

$$f(\theta) = f(\theta_1 + \theta_2) = \begin{cases} \omega & , \text{if } 85^\circ \leq \theta \leq 95^\circ \\ 0 & , \text{otherwise.} \end{cases}$$

During the initial phases of evolution, when the neural networks are mostly untrained, all networks may direct the *Mekahand* to grasp at positions where it cannot even touch the object. As a result, in early generation f_1 is often effectively zero. Thus in this stage, f_2 , which rewards approaching the target object, is important for differentiating the fitness. After further evolution, when the hand can grasp the object, f_1 begins to dominate and the neural networks are ranked mostly by grasp quality. In addition, the third term f_3 is large if the *Mekahand* is not blocked by obstacles (e.g. objects other than the target object). Finally, the fourth term (f_4) rewards facing the palm of the robotic hand towards the target object. Parameters $\alpha, \beta, \gamma, \lambda, \epsilon$ and ω adjust the relative effects of those four terms. In this way, the described fitness function rewards ANNs first to learn to approach the object, and then to grasp the object in an increasingly appropriate way. Algorithm 1 shows the fitness function in detail.

5.2.4 Reducing Training Time through Parallelization

The computational cost incurred by the *sequential* implementation of the fitness function computation is as follows. For one experiment, each generation consists of \hat{o} ANNs, and each ANN is evaluated over \hat{s} object combinations. Each object combination contains \hat{b} objects, and each object is selected as \hat{k} candidates to be an input. If one experiment runs for \hat{g} generations, the total number of independent training simulations in GraspIt! T is $\hat{o} \times \hat{s} \times \hat{b} \times \hat{k} \times \hat{g}$. In our experiments, $\hat{o} = 200, \hat{s} = 5, \hat{b} = 4, \hat{k} = 10, \hat{g} = 150$. Thus, $T = 6,000,000$. Therefore, a parallel strategy that dispatches different trials to all available computer cores is implemented to encourage computational efficiency. In particular, work is dispatched over a network to multiple GraspIt! processes that run on different computers. In this way, each CPU core in different computers can be fully employed, and the resulting multi-threaded implementation speeds up the evolution process.

Alg 1 Computation of the Fitness Function

1: **Input:** Q is the grasp quality after the execution of a single grasp, θ is the summation of θ_1 and θ_2 , \vec{P}_i is the predicted position of hand for grasping by the network, \vec{O}_i is the coordinate of the selected object after the mouse click, \vec{S}_i is the actual hand coordinate after interacting with the environment.

2: **Output:** A fitness evaluation of a single grasp.

3: Let A_j be a set of 3D coordinates of objects in the environment, where $1 \leq j \leq n$;

4: **for** $j = 1$ to n **do**

5: $Dist_o = \min(Dist_o, \sqrt{\sum_{i \in x,y,z} (\vec{A}_{j,i} - \vec{S}_i)^2})$;

6: **end for**

7: $Dist_t = \sqrt{\sum_{i \in x,y,z} (\vec{O}_i - \vec{S}_i)^2}$;

8: **if** $(Q = 0) \parallel (Dist_o < Dist_t)$ **then**

9: {*No grasp quality or *Mekahand* is closer to other objects.*}

10: $f_1 = 0$;

11: $f_2 = \frac{\beta}{d(\vec{P}_i, \vec{O}_i) + \alpha}$;

12: **else**

13: $f_1 = \gamma Q$, where $\gamma \geq 10000$;

14: $f_2 = k$, where $k \leq 1000$;

15: **end if**

16: $f_3 = \frac{\lambda}{d(\vec{P}_i, \vec{S}_i) + \epsilon}$;

17: **if** $(85 \leq \theta) \ \&\& \ (\theta \leq 95)$ **then**

18: **if** $Dist_t < 50$ **then**

19: $f_4 = \omega$;

20: **else**

21: $f_4 = w$, $w < \omega$;

22: **end if**

23: **else**

24: $f_4 = 0$;

25: **end if**

26: **return** $sum = \sum_{i=1}^4 fit_i$;

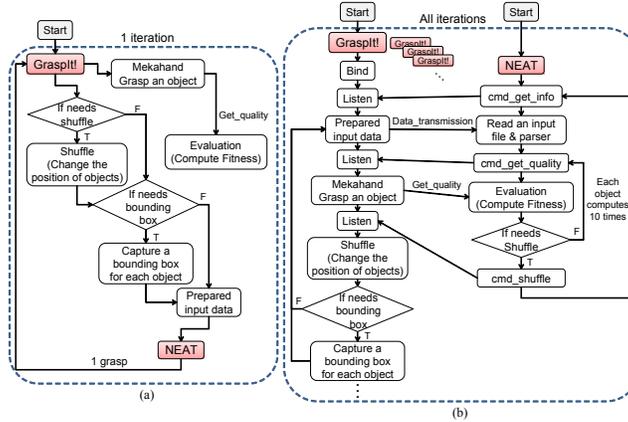


Figure 8: The same computers were used to compare the sequential and parallel comparison methods. (a) The original sequential method. (b) The faster parallel method. The results show that with the original sequential implementation, the program only utilizes a single core, but after parallelizing the algorithm, the program can fully utilize four cores, and the experiment’s run time is shortened by a factor of three.

Figure 8(a) illustrates the sequential method for each generation. To reduce execution time, the following computational steps were parallelized, as shown in Figure 8(b). The sequential method was redesigned as a producer-consumer pattern for parallel programming, which consists of two major components: GraspIt! simulator and NEAT training. NEAT as producer generates the postures and enqueues into the queue for further processing. GraspIt! simulator as consumer dequeues the postures and run the evaluation of grasping quality, and then enqueues the quality measurements into the second queue. After all postures were generated and NEAT obtained the grasping qualities, the information would be fed into the neural network as an input for the next generation. First, three commands are defined: *cmd_get_info* is to get the depth array, *cmd_get_quality* is to get the quality for each grasp, and *cmd_shuffle* is to change the position and orientation of each object. Here, assume that four instances of GraspIt! are run and waiting for commands. Two kinds of threads are created: *Organism tasks* that use ANNs from NEAT’s main process to generate grasping tasks and collect the resulting fitness score; and *GraspIt! tasks* that communicate with a GraspIt! process to send the output from an ANN for simulation in GraspIt!, and receive the resulting grasp quality. The speedup achieved by such parallization depends on how many GraspIt! Instances are running. To gain more computing power and speed up the training time, we implement the producer-consumer as client/server architecture where command and data are exchanged. Our results show that the run time is accelerated by at least a factor of three.

5.3 Robotic Control in the Imprecise Computation

Our specific robotic grasping task has time constraint, so how to maximize the reward associated with the optional part of execution while satisfying all mandatory deadlines, can be considered it as an imprecise scheduling problem. For example, objects localization from images processing, rough estimate of location from low-resolution images produces in time, whereas accurate location from high-resolution images takes longer time. In the robotic control, the mandatory part guarantees an approximate solution, so can be viewed as mandatory. On the other hand, the optional part depends on the precision of the solution and must be non-decreasing, so can be varied by adjusting robotic arm speed, neural network training time. Ideally, each task would finish running its optional

part, but these computations can be canceled when out of time. Therefore, this model is still not sufficient to provide a well-defined scheduling problem, because some mechanisms are required to decide which optional parts to carry out. This paper fulfills the fundamental step to explore the boundary region of tolerance and find best effort techniques to satisfy the minimum QoS requirement. With an increased number of subsequent deadline constraints, the objective is to derive the relationship between execution time and grasping trajectory accuracy. The tradeoffs analysis lays the foundation for the scheduling in the imprecise computation framework. For example, when we remotely operate *Dreamer* in the real world, the choices of completion time associated with each data point might be different in the experiments, as described in Section 7. In this case, the tradeoff mapping could be of an aid to assist the robot to make a decision on choosing different combinations in order to complete the task without exceeding the time constraint.

Next, we shall present the experimental quality evaluation.

6 EXPERIMENTAL EVALUATION

In this section, the training and testing experiments are described. The design and parameters are presented in 6.1 and 6.2. The first set of training experiments combines the four fitness components in different ways, as described in 6.3. The best combination is applied in the second and third sets of training experiments, which evaluate the benefit of applying a bounding box to focus the ANN's attention in 6.4. Fully trained ANNs are tested in simulation (6.5) and also transferred to the real robot (6.6).

6.1 Experimental Design

Because the raw depth data from the Kinect sensor is of high dimensionality, for practical purposes the array is first down-scaled. Before the input data is supplied to an ANN, the 640×480 pixel array was sampled to form a reduced 20×15 array. A larger scale was also tried, such as 40×30 , 80×60 , but not only the evolution process was so time-consuming but also the improvement over the results was not obvious at all. Therefore, we decided to shrink back 20×15 array. This smaller array was converted to gray-scale intensity values, and then normalized between zero and one; an example is shown in Figure 9. The input data also includes a coordinate that represents the mouse click input from the user that specifies the target object. In the grasping experiments, the coordinate is chosen by randomly picking a different point on the target object in each trial. To increase accuracy in evaluating each network, they are each evaluated five times over different trials. That is, the robot attempts to grasp each target object five times, and the fitness value is the average over all the attempts. To preserve generality, the position and orientation of the objects for each evaluation are randomized.

The experiments are divided into two parts: training and testing. A collection of objects are divided into N separate classes, and for each class, ANNs are trained by NEAT to grasp objects from that class. For testing, the best neural network generated from training is further tested in simulations over objects placed in different locations. A final test applies a real scenario from *Dreamer* to the evolved neural networks. The flowchart for training and testing is shown in Figure 10. All experimental parameters are described in 6.2.

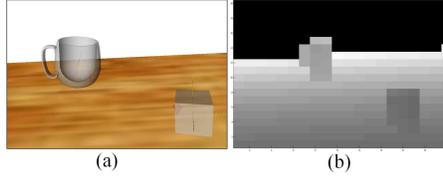


Figure 9: Sample input data for training neural networks. (a) The RGB pixel data of the scene from the camera within GraspIt!. (b) The 20×15 depth data array supplied to the neural network as input. The depth data is normalized to a floating point number between $[0, 1]$. The purpose is that the original raw pixel data is high-dimensional, so a down-scaled data of the same data can be easily performed in practice.

6.2 Experimental Parameters of Neural Network

In the experiments, the population size was set to 150–200. Different values of the three parameters α , β and γ of the fitness function (Eq. 16) were tried and tuned to guide evolution. The number of generations was 100. The coefficients for measuring compatibility for NEAT were $c_1 = 1.0$, $c_2 = 2.0$, $c_3 = 2.0$. The survival threshold was set to 0.2–0.3. The drop-off age was set to 10–20. Recurrent connections were disabled because the task is not dependent on history. The probability of adding nodes and adding new connections to evolved ANNs were set to 0.2 and 0.3, respectively. Detailed description of these parameters are given in [51].

6.3 Testing Combinations of Fitness Function Components

Training experiments are performed with four target objects plus a dining table to vary the scene distribution as shown in Figure 11(a), (b). The goal is to gauge which combination of fitness function components (from Section 5.2.3) will yield the best performance. Figure 11(c) shows the five results of the fitness function for four scenarios through iterative training experiments.

Because grasp quality (f_1) is the most important performance metric, each case must contain f_1 , so the combination of total cases is $C_3^3 + C_2^3 + C_1^3 = 7$. The following notation is used to refer to the section: $\{ \overline{Fit}_i \mid i \in \{12, 13, 14, 123, 124, 134, 1234\} \}$.

As an example, \overline{Fit}_{134} denotes the case with f_1 , f_3 and f_4 . The simulation environment performs a series of simulated grasps on one object on a dinner table for grasping evaluation. Figures 12(a)-(d) show training results for grasping a single cylinder, cube, sphere, and mug, respectively. For the cylinder, cube, and mug, the maximum grasping quality f_1 is achieved through \overline{Fit}_{1234} (i.e. each fitness component is helpful). However, Figure 12(c) shows that the maximum grasping quality f_1 for a sphere is achieved through \overline{Fit}_{134} , which suggests that f_2 does not contribute to better performance. Because the sphere is relatively small, placed in-between other objects it is sometimes blocked by other objects. Because its color is similar to the table, it is hard to distinguish it from the other objects. As a result, NEAT will be misled by the simple f_2 distance metric.

6.4 Bounding Box Experiments

In this section, the approach is applied to a human-supplied bounding box to focus the robot’s visual processing on the target object, thereby reducing its dimensionality.

In the second set of experiments, four different training scenarios (without a bounding box) are performed with different target objects, similar to Section 6.3. Figure 13(a)-(d) show training results for networks trained to grasp a single cylinder, a single cube, a single sphere and a single

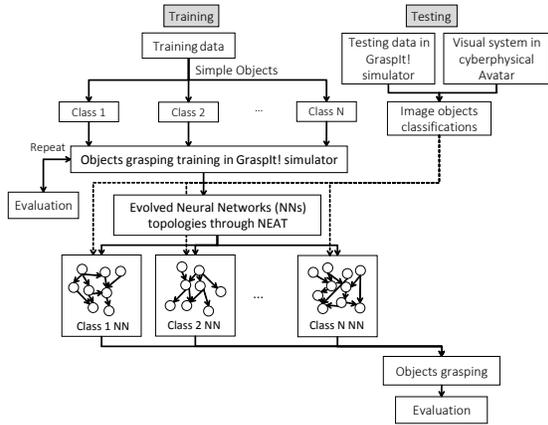


Figure 10: The flowchart of the training process and the testing process for the experiments. In the training process, a set of objects are grouped into N separate classes, and then each class produces a neural network through NEAT; in the testing process, the best neural networks can be applied in simulations and tested in a real scenario. The grasping accuracy can be further improved by preprocessing the data before conducting training/testing experiments. These processes can examine if the proposed approach can work.

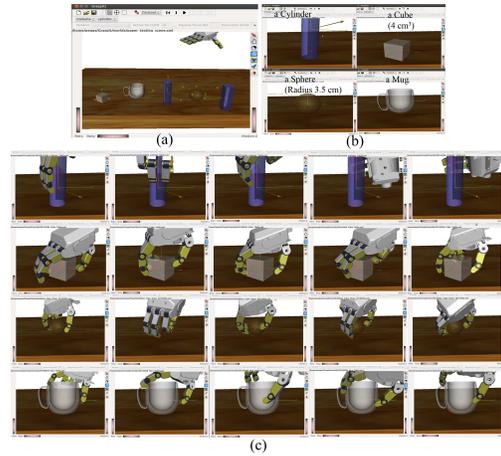


Figure 11: Experimental scenarios. (a) A single cylinder, cube, sphere, mug, and cuboid with a dining table and the *Mekahand*. (b) Focus on a single target object each time. (c) The five results for each object during training. The conclusion is that the fitness function can guide *Mekahand* to grasp four different objects.

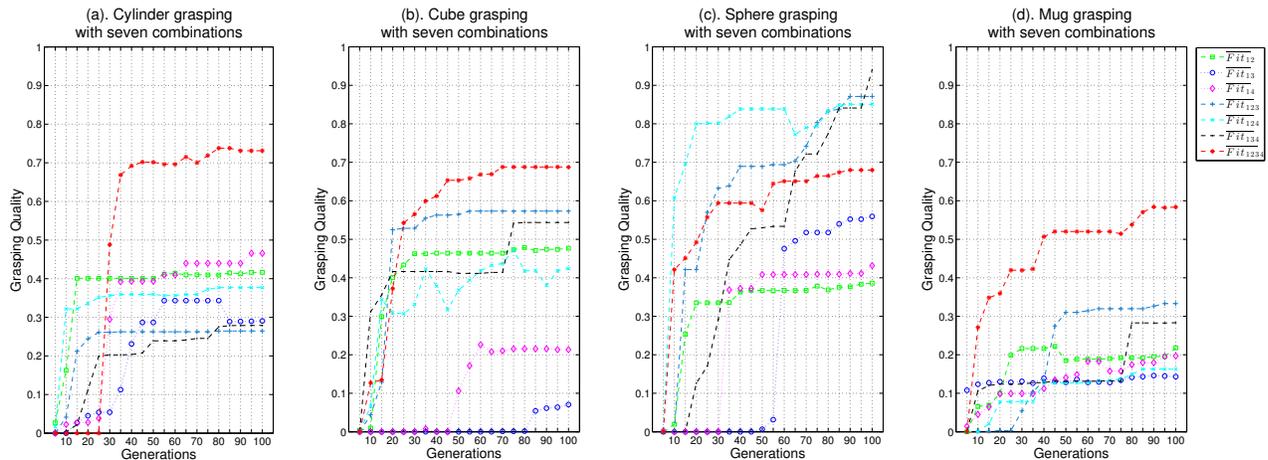


Figure 12: Training performance with combinations of fitness components. The training scenario includes a cylinder, a cube, a sphere, and a mug, on a dinner table, but the depth sensor focuses only on a single object for each experiment. The x axis represents the number of generations while the y axis represents the normalized grasping quality. These figures show how grasping quality increases over the course of evolution. To evaluate whether each of the four fitness component helps improve performance, (a)-(d) compare seven combinations of fitness components: (a) shows the results for grasping the cylinder, (b) for the cube, (c) for the sphere, and (d) for the mug. The conclusion is that Fit_{1234} produces the best grasping quality for (a), (b), (d), while Fit_{134} provides the best one for (c).

mug. These figures show how fitness values increase over the course of evolution. Note that larger fitness value implies better grasping quality; also, to differentiate the contributions of f_1, f_2, f_3 and f_4 , each of these terms is normalized.

According to the best combination of the four fitness components from Figure 12(a)-(d), Figures 13(a), (b), (d) differentiate the contributions of f_1, f_2, f_3 , and f_4 , and Figure 13(c) differentiates the contributions of f_1, f_3 and f_4 . Note that the maximum score f_1 can attain is 0.6, the maximum for both f_2 and f_3 is 0.1, and the maximum for f_4 is 0.2. Because f_2 and f_3 encourage approaching objects and avoiding obstacles, and f_4 rewards orienting the palm toward objects, that can serve as secondary objectives. These terms are therefore given lower weights than f_1 , which measures the grasping quality itself and is thus the most important performance metric.

Because in practice only the best controller would be used, overall best-case results are presented here. To start evolution, individuals in the population are initialized with random weights and a simple topology (i.e. input nodes fully connected to one hidden node, and this hidden node fully connected to the outputs). Because randomly generated policies generally do not cause the robot hand to approach the target objects, low fitness scores are expected. In this stage, f_1 for all the networks is low, so the fitness scores of the networks are mainly determined by f_2 and f_3 . These two terms guide evolution to produce networks that approach the objects without being blocked by obstacles. The f_4 component leads the *Mekahand* to the right orientation toward the object. In accordance with this explanation, Figure 13(a) shows that initially f_1 is smaller than f_2 and f_3 . However, after 25 generations, f_1 becomes dominant. Then, after 90 generations, f_1 reaches its maximum value of 0.5, which means the *Mekahand* can grasp the object more accurately with proper position and orientation. Similar results appear in the other three experiments (Figure 13(b)-(d)). In Figure 13(b), after approximately 15 generations, f_1 sharply increases, and the total fitness value steadily increases to reach a maximum value of 0.7. In Figure 13(c), only f_1, f_3, f_4 are considered, but the fitness value remains around 0.7. In Figure 13(d), the fitness value only achieves 0.6. The reason is that it is difficult for the neural network to distinguish the mug object from the other objects. Comparing the four figures, it can be seen that the fitness scores of neural networks trained on the simple objects (Figure 13(a)-(c)) were larger than those trained on the more complicated one (Figure 13(d)). However, even in the more complicated scenario the networks all learned to approach the target objects and grasp them.

The third set of experiments tests evolution in the same four scenarios, but adds a visual bounding box that can focus the ANN on the most relevant information. The first experiment is shown in Figure 13(e). The fitness value gradually increases, and after 5 generations, the values are better than Figure 13(a), achieving a value of 0.9 after 45 generations. Similar results are seen in Figure 13(f)-(h). In Figure 13(f), the maximum fitness value is 0.92. Figure 13(h) illustrates that with a bounding box, more complex object configurations can still produce consistent results around 0.8.

This experiments suggest that the more complex the training scenario (i.e. the number of different kinds of objects in the scene), the more difficult it is to train the neural network. Furthermore, if a facet is obscured or the depth array values of an object are similar to the background, then even if the object to be grasped is simple, the training results are poor. However, applying the bounding box significantly improves the results in such cases.

6.5 Validating the Generality of Evolved Neural Networks

The training methodology results in neural networks evolved to grasp objects in simulation. To validate such networks, they were further tested in a variety of novel situations through GraspIt! (i.e. situating for which object was not explicitly trained). Most objects in the scenes were not

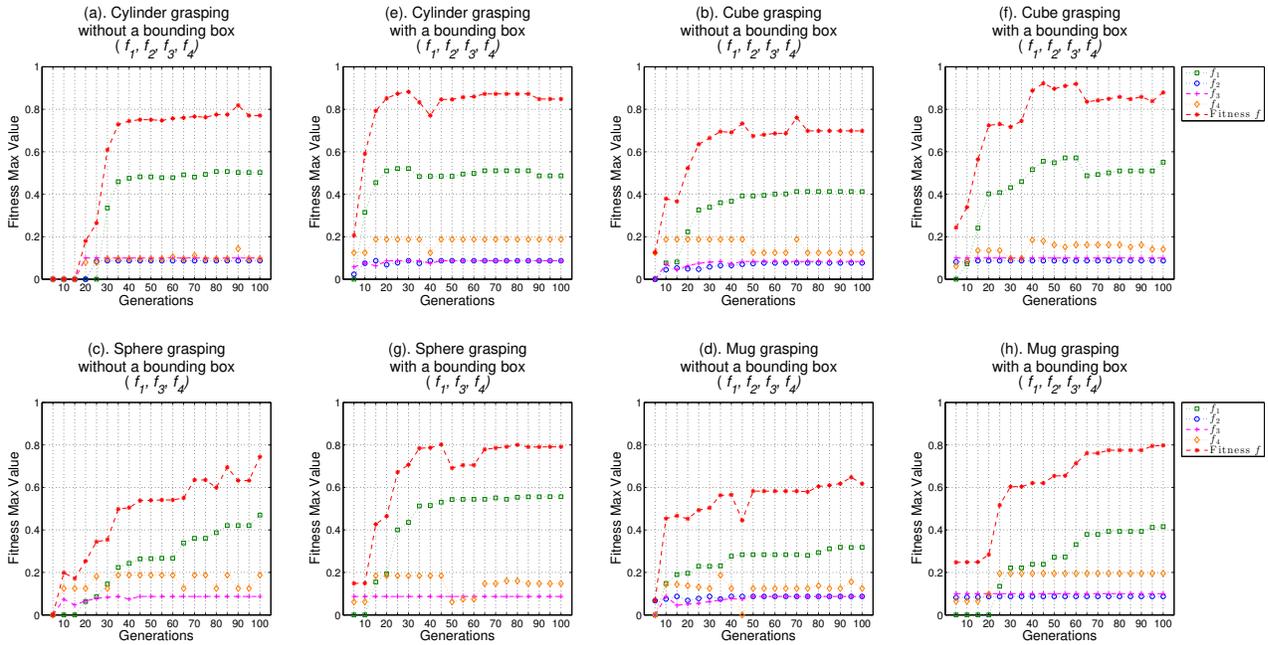


Figure 13: Training performance with and without a bounding box. How fitness values increase over generations is shown for each experiment. Plots (a) and (e) show a scenario with a single cylinder on a table, (b) and (f) a single cube on a table, (c) and (g) a single sphere on a table, (d) and (h) a single mug on a table. To evaluate whether a bounding box benefits performance, (a)-(d) have no bounding box, while (e)-(h) include the bounding box technique. The total fitness value is shown, as are the contributions from the three or four underlying normalized terms. The conclusion is that the bounding box increases performance, and all experiments eventually evolve ANNs able to grasp the objects in simulation.

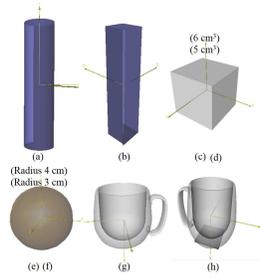


Figure 14: Testing different sizes and textures of objects across novel locations and orientation. Shown in the figure are a cylinder, a cuboid, a cube, a sphere, a mug and a plated mug. Note that the letters labeling each object correspond to similar labels in Table 1.

#	Objects	Class	without a Bounding Box	with a Bounding Box
(a)	cylinder	N_A	52%	89%
(b)	cuboid	N_A	65%	81%
(c)	cube (6 cm^3)	N_B	69%	76%
(d)	cube (5 cm^3)	N_B	73%	82%
(e)	sphere (radius 4 cm)	N_C	71%	88%
(f)	sphere (radius 3 cm)	N_C	68%	80%
(g)	mug	N_D	71%	85%
(h)	plated mug	N_D	62%	74%
	Mean/Std		66.38% ($\pm 6.80\%$)	81.88% ($\pm 5.33\%$)

Table 1: Generalization results of grasping objects at novel positions with evolved networks. The results with a bounding box outperform the ones without a bounding box, which indicates that a bounding box is an effective way of increasing grasping performance.

seen at all during evolution or not placed in the same location, and their arrangement is new. The experiment thus measures how general the evolved solutions are. A successful case is recorded if the *Mekahand* can grasp the object; otherwise it is recorded as a failure.

For this generality test, each object was tested 100 times. The grasping procedures were implemented under test conditions randomly placing the different sizes and textures of a cylinder, a cube, a sphere, and a mug, as shown in Figure 14, at different positions and orientations on the table. The evolved neural networks in 6.4 were labeled as Cylinder (N_A), Cube (N_B), Sphere (N_C), and Mug (N_D), and based on similar classification of objects, the most appropriate neural network was chosen for testing. The success rate in Table 1 compares the neural networks with the different objects. These results show that despite its simplicity, the proposed bounding box method still performs reasonably well in grasping novel objects. However, if the target object is too far from the center of the image frame, the neural networks often perform unreliably, indicating the training process may need further refinement to deal with such boundary cases. Table 1 shows the best results from among all the experiments. Also, in some cases the *Mekahand* collides with objects while grasping, because many objects are placed on the table. A potential remedy is to decompose the movement into more steps to avoid such collisions. One way to do so would be to rely on additional input from the human supervisor.

6.6 Validating with Dreamer

Beyond simulated results, learned policies were also transferred to the physical world. A physical (i.e. not simulated) Kinect sensor was applied to capture object depth array information. This information was provided as input to an evolved neural network to guide *Dreamer* robot’s grasp.

6.6.1 Kinect Sensor Implementation

To retrieve the Kinect sensor data and feed it into the system, the sample program regview provided by OpenKinect project¹ was modified. This program was enhanced to be run as a server that waits for the connection from the remote-control PC over the TCP/IP. Besides, it was tweaked to register

¹<http://openkinect.org>

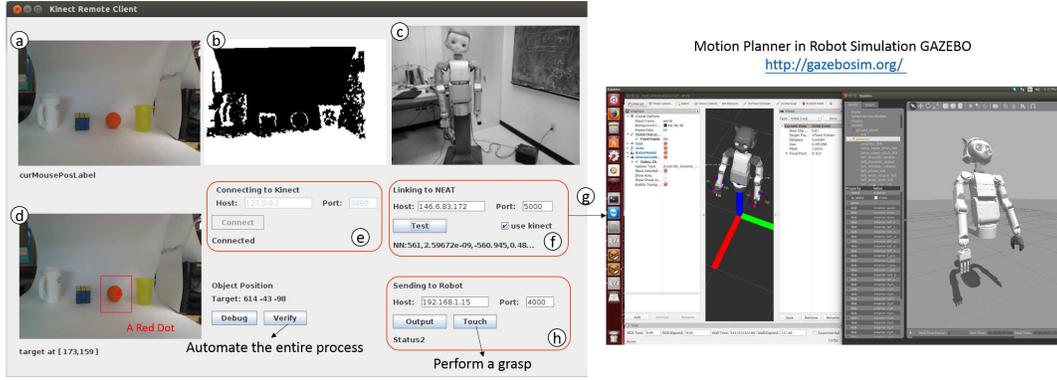


Figure 15: A screen capture of the remote-control software application for supervising the *Dreamer* robot. (a) Color and (b) depth images from the Kinect sensor. (c) The image from the IP camera. (d) An image snapshot taken when the user clicks on the color image. (e) A dialog for connecting to *Dreamer* through a computer network. (f) A dialog for inputting the captured depth array into an evolved ANN. (g) Use motion planner to obtain a trajectory. (h) A dialog for sending the orientations and positions from the ANN to *Dreamer* to control its grasp. The conclusion is that the grasping experiment can be implemented through the remote control panel.

the video format as `FREENECT_DEPTH_REGISTERED`. The reason is that in the Kinect sensor, the depth camera and the color camera are two separate sensors, which means their views are different. Only by doing so, the depth data will be projected to the view of the color camera. In this video mode, the depth data is in millimeters, and the pixel coordinates can be translated from (i, j, z) to (x, y, z) as follows:

$$x = (i - \text{width}/2) * (z + \text{minDistance}) * \text{scaleFactor} * (\text{width}/\text{height}), \quad (17)$$

$$y = (j - \text{height}/2) * (z + \text{minDistance}) * \text{scaleFactor}, \quad (18)$$

$$z = z, \quad (19)$$

where weight and height are the images size. The x, y, z is a right-handed Cartesian system: with z axis perpendicular to the Kinect image towards the image, x -axis points to the left, and y -axis points up. Before sending commands to the robot, the coordinates are transformed again to match *Dreamer*'s coordinates.

6.6.2 Remote Control Panel

Figure 15 shows a screen capture of the remote-control application for supervising *Dreamer*. The remote-control user interface shows the color images and depth images from the Kinect camera; the images from the IP camera are displayed in the third image panel at the top of the user interface.

To automate the high-level supervision of the grasping experiment, six commands was implemented on the remote control panel: Connect, Test, Output, Touch, Verify, and Debug, as shown in Figure 15. When the supervisor clicks on the Connect button, the computer connects to the Kinect sensor to capture depth information. Then, when the Test button is pressed, the depth array is provided to the evolved neural network as an input. After executing the neural network, its outputs are interpreted as coordinates and orientation of the hand for grasping the object. When the Output button is pressed, the results are sent to *Dreamer* and the robot is directed to approach the object. Finally, when the Touch button is pressed, *Dreamer* will grasp the object using the grasping

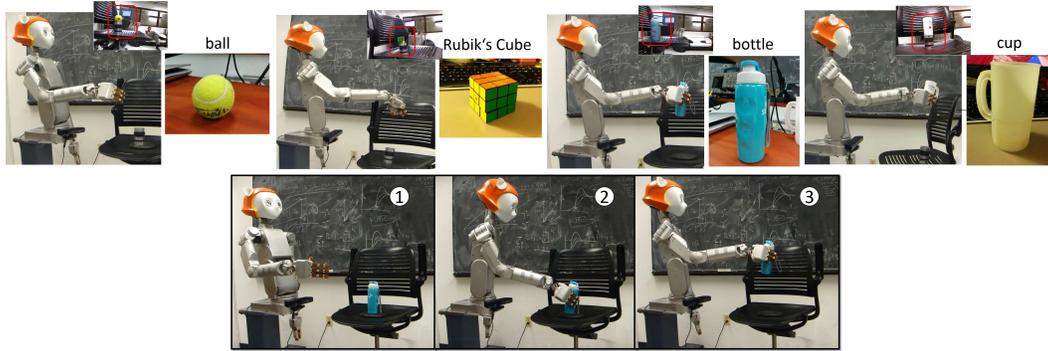


Figure 16: Screen captures from videos (<http://www.cs.utexas.edu/~peggy/rtss2015.html>) demonstrating *Dreamer* grasping a ball, a bottle, a cube and a cup through an evolved controller. Note that the small picture with red dots are the snapshots from Kinect sensor panel. The bottom snapshots labeled with (1)-(3) represent the object grasping process, from the initial, approach to grasp a bottle. The figures confirm that transferring results from simulation to reality is possible, and applying the approach generalizes to novel objects.

information provided by the neural network. After *Dreamer* obtains this information, i.e. the grip orientation and position, the controller PC computes the distance between the *Mekahand* and the object, predicts the hand's trajectory, and approaches the object. Once the *Mekahand* is near the target object, the thumb and the three finger motors are synchronized to perform the grasp. A Verify button is provided to automate the entire process for convenience; the Debug button serves to aid in system debugging, providing coding logging information.

6.6.3 Transitioning to Physical Controller

An automated grasping platform was built to demonstrate this process. The networks evolved in simulation are transferred to this platform to evaluate them in a physical environment. To carry out an experiment, a human experimenter uses the control panel to choose a target either without or with a bounding box in the color image from the laptop screen with the Kinect sensor by clicking on it. After designating the target, a copy of the color image is copied to the target object panel, and a red dot is added on the image, indicating the position of the click. The depth data at that point is used to calculate the approximate position of the object to be grasped. This results specifies the grasping task for the robot to perform. Note that the grasping behavior was not evolved on the actual robot, but was transferred from simulation.

The video <http://www.cs.utexas.edu/~peggy/rtss2015.html> demonstrates grasping of novel objects from the simulation to the real *Dreamer* robot. In addition, Figure 16 shows screen captures taken from a proof-of-concept demonstration of grasping a tennis ball, a bottle, ball, a Rubik's cube and a cup. *Dreamer* can successfully approach and grasp target objects when controlled by an evolved neural network.

Since these objects were not seen during evolution, the experiment demonstrates two achievements: (1) learning transfers from simulation to the real world, and (2) it generalizes to grasp objects. Quantifying how well grasping works needs a metric for the assessment of the quality of a real grasp, so further work to incorporate real sensor data on the *Mekahand* (e.g. touch pressure) is ongoing.

7 THE REAL-TIME PERFORMANCE MEASUREMENTS AMONG TASK COMPLETION TIME, TRAINING EFFORT AND GRASPING QUALITY

This section discusses whether string task completion deadlines can be met by applying imprecise computation to trade increased speed for decreased accuracy. That is, in some situations a faster yet less precise grasp may better satisfy the use case. When considering tradeoffs, the design space of the grasping task has three main dimensions: (1) the training effort, (2) the task completion time, and (3) the grasp quality [12]. The next sections explore tradeoffs between these dimensions. In 7.1, we first investigate how increasingly stringent time limits on computation reduce the accuracy of the robotic hand’s approach trajectory. We then discuss the performance tradeoff between grasp quality and task completion time. In 7.2, we evaluate the tradeoff between training effort and grasp quality, and in 7.3, we measure the tradeoff between training effort and task completion time with the successful grasp.

7.1 Grasp Quality vs. Task Completion Time Tradeoff Evaluation

For the grasping task, trajectories may be denoted by the point-to-point positions and orientations of the end-effector as long as no collision occurs. This section focuses on the actual interaction between the *Mekahand* and its environment assuming that there is no collision.

In a grasping experiment, the initial starting point S^* is the current position and orientation of *Dreamer’s* end effector. The human supervisor assigns an object to be grasped from the user interface panel; the evolved neural network automatically determines the final destination D^* and orientation of *Mekahand* and sends it to *Dreamer’s* main controller. On command, *Dreamer* moves along the designated trajectory to approach and grasp the object, and then returns to the start position S^* . The actual trajectory of *Dreamer* is acquired by recording the position of the end-effector from forward kinematics calculations with joint positions. In controlling the movement of *Dreamer’s* arm, we use the proportional-derivative (PD) controller in the Whole Body Control (WBC) algorithm. The position and orientation data are transmitted with a wireless system from the sensors to the control computer.

The first set of experiments measures the quality of the grasping trajectories versus various task completion times. In each experiment, *Mekahand* moves from S^* to D^* within a specified time interval of length ranging from 8 seconds down to 0.5 second. Each configuration was measured five times over different trials to obtain an accurate trajectory error estimation. An ideal trajectory was designed by a trajectory generation algorithm, and all experiments attempted to follow this trajectory, subject to different completion time deadlines. Each execution time was separately conducted five times, and averaged the five trajectories, as shown in Figure 17. In Figure 17, the highest variance in error are found in the 0.5 second trials, while the lowest variance in error are found in the 8 second trials. The differences between the ideal and actual trajectories for a one-way trip were recorded and depicted as boxplots for different time constraints and trials. Each scenario (execution time) was performed five times, and nine scenarios of experiments with the various execution times from 0.5 second to 8 seconds were tested. Therefore, forty-five results were yielded, as depicted in Figure 17 (a)-(i). Figure 17 illustrates that the five trials for each execution time are comparable in the trajectory distributions. Figure 18 shows the difference between the ideal and actual trajectory over various task completion times. As expected, the trajectory closest to the ideal one is the one given the most time (i.e. 8 seconds). Figure 18 shows that in general,

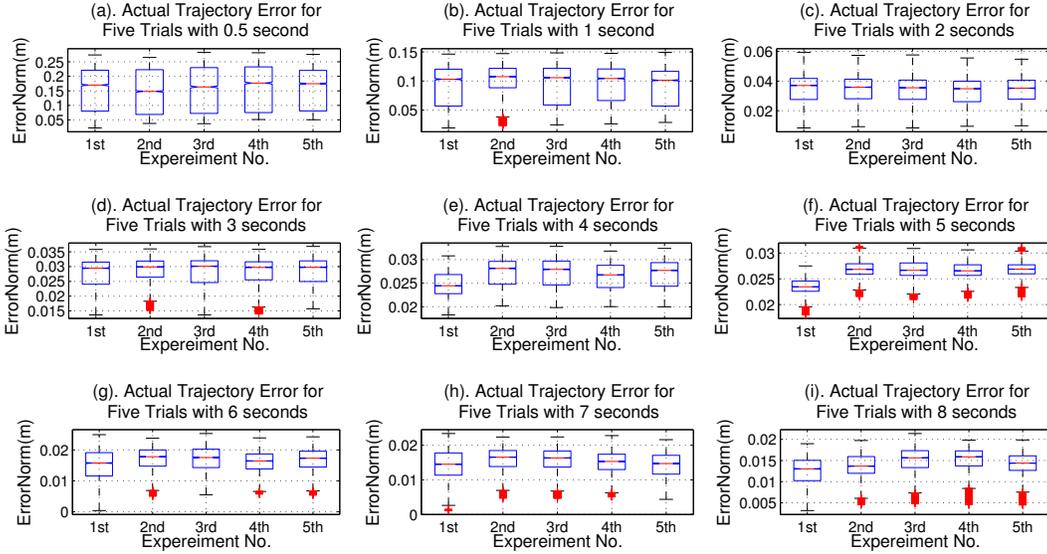


Figure 17: The results at five trials for nine scenarios with different execution times. The x axis indicates the trial number; the y axis indicates the normalized trajectory error compared to the ideal trajectory across the entire trajectory. Figures(a)-(i) show trials with execution times ranging from 0.5 to 8 seconds, summarizing in total the distribution of trajectory errors for 45 trials. The conclusion is that trajectory distributions for trials of particular length are similar enough to justify deriving statistics models.

the shorter the completion time deadline, the higher the trajectory error. It should be noted that the design goal is to contain the trajectory error so that the grasp action can succeed at the end of the trajectory.

To predict the probability of a success grasp, we can fit a statistical model to characterize the tradeoff between average trajectory error and task completion time. First, it is necessary to evaluate whether the five experiments are sufficient to represent ground truth. The standard deviation (STD) (root-mean-squared (RMS)) error for each task completion time respectively range from 0.00231(m) to 0.00009(m) (from 0.15742(m) to 0.01395(m)). Such a low STD indicates that the results do not vary much, and can therefore serve as a reasonable basis to derive models of RMS error.

To find a well-fitting regression model, Table 2 shows the results of approximated RMS error and adjusted R^2 of power, Weibull, rational, Gaussian and polynomial distributions. The trajectory errors are best modeled by a polynomial distribution of order five, as demonstrated by its lowest RMS errors and highest coefficient of determination. The best fit polynomial model function $f_{\text{model}}(x)$ is:

$$\begin{aligned}
 f_{\text{model}}(x) = & -0.000714 * x^5 + 0.002011 * x^4 \\
 & + -0.02155 * x^3 + 0.1088 * x^2 \\
 & + -0.26 * x + 0.2635.
 \end{aligned} \tag{20}$$

Figure 19 shows the RMS errors fit to a linear interpolation and polynomial distributions. This model can predict the error after 9 seconds that is nearly stable (below 0.02 m) which indicates the 9 results are sufficient to proceed with the following experiment. To investigate how the reduction of task completion time may jeopardize sufficient accuracy for effecting a grasp, 6 seconds was chosen as the time-constraint for the following experiment.

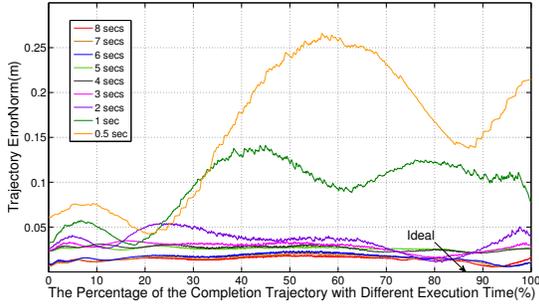


Figure 18: Tracking trajectories by varying execution times ranging from 8 to 0.5 seconds. The x axis represents the completion of the trajectory while the y axis represents the normalized trajectory error compared to the ideal trajectory. The highest error is found in the 0.5 second trials, while the lowest error is found in the 8 seconds trials. The trajectory error increases as the allowed time for execution decreases.

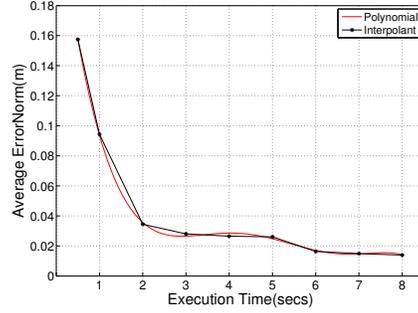


Figure 19: Fitting the RMS data with linear interpolant and 5-degree polynomials. This chart shows that inaccuracy is maximal when execution time is shortest (0.5 seconds), but rapidly improves as the budget increases to 3 seconds. Error decreases slightly between 3 and 6 seconds, and plateaus thereafter. The derivation line shows that the results approaches stability after 6 seconds, even the experiments after 9 seconds still can predict the error may be below 0.02 m.

The second set of experiments focuses on trajectory accuracy and latency delay within a given time constraint, i.e. six seconds. The controller operation is the most time-consuming part of the practice, because many DOFs are considered in each step. In contrast, sending data and NEAT operation are relatively fast, only taking around 1-2 seconds each. Therefore, a round-trip is given fourteen seconds. The latency delays are 100ms, 50ms, and 10ms. The longer the latency delay, the worse the performance. In order to complete the task within the limited time frame, it is hard to control *Dreamer* very well. Figures 20(a)-(i) depict the relationship between the position and time with 100ms, 50ms, and 10ms; Figures 20(j)-(l) display the relationship between error and time. Note that the error was computed as the actual minus the ideal desired trajectory. Figure 20, as expected, shows that the lowest latency delay (10ms) performs best.

7.2 Training Effort vs. Grasp Quality Tradeoff Evaluation

We first describe the experimental setup and then present a set of grasping results that relate the quality of grasping to the training effort (defined to be the time spent on searching for the best *Mekahand* configuration for effecting the grasp by the NEAT algorithm). To speed up the training computation, we apply a parallelization strategy and run the NEAT algorithm with four multi-core computers.

We evaluate the effectiveness of our learning approach by conducting the two following sets of experiments. For the first experiment, the computational cost incurred by the *sequential* implementation is described in 5.2.4. The *parallel* strategy which dispatches different trials to all available computer cores is implemented to increase computational efficiency. In particular, work is dispatched over the network to multiple GraspIt! processes (thirty-six threads) which run on four computers, whose specifications are detailed in Table 3.

The first set of simulated experiments involve scenarios with different target objects, as explained in Section 6.3. Specifically, Figure 21(a)-(d) show the training results for networks trained

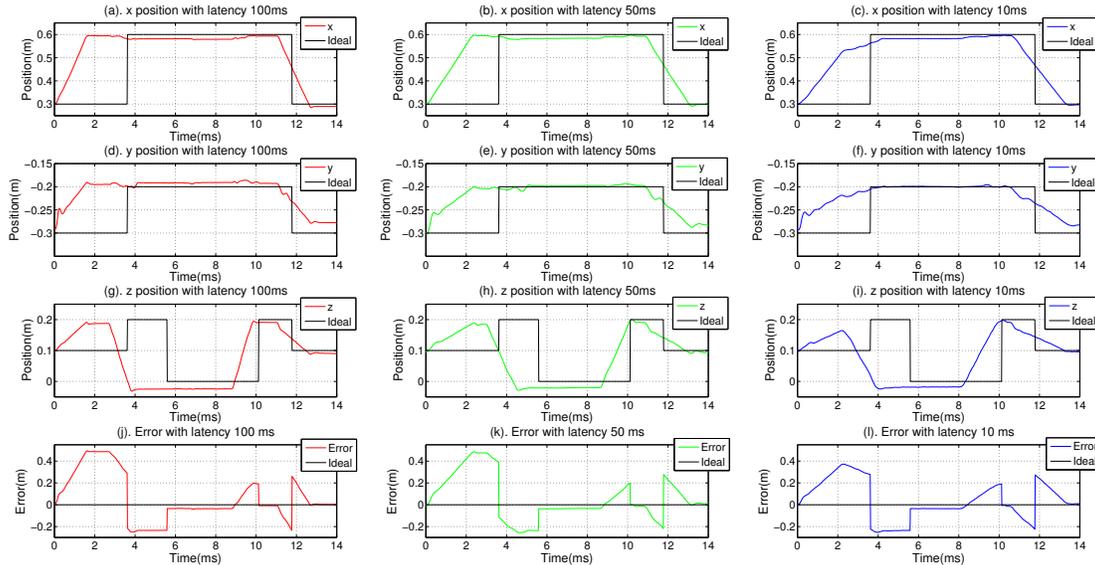


Figure 20: (a)(d)(g), (b)(e)(h), and (c)(f)(i) display the relationship between trajectory and latency delay (100ms, 50ms, and 10ms); (a)-(c), (d)-(f), and (g)-(i) show x , y , and z positions in 3D space. Figures(j)-(l) depict the relationship between the error and latency delay. The x axis represents the given time; the y axis in (a)-(i) denotes the actual and ideal positions, and in (j)-(l) denotes the error. Here, the error was calculated as the actual minus the ideal trajectory. As expected, the variation with 10ms was the best.

Fitting model	RMS errors	adjusted R^2
Power (2 terms)	0.006670	0.9812
Weibull	0.006179	0.9838
Rational (degree 5)	0.003164	0.9958
Gaussian (2 terms)	0.002755	0.9968
Polynomial (degree 5)	0.002308	0.9977

Table 2: A comparison is shown of five well-fitting statistical models of how RMS error varies with execution time budgets. The best fitting model is a polynomial of degree five, which is used to predict the expected RMS error at 9 seconds.

CPU	GPU (NVIDIA GeForce)	Memory (DDR3)
Intel Core i7 @ 3.6 GHz 4 Core / 8 Threads	GTX 760	16GB @ 1333 MHz
Intel Core i7@ 3.0 GHz 8 Core / 16 Threads	GTX 980	32GB @ 2133 MHz
Intel Xeon@ 2.67 GHz 4 Core / 8 Threads	GTX 285	8GB @ 1066 MHz
Intel Core i7@ 2.8 GHz 2 Core / 4 Threads	GT 520M	8GB @ 1333 MHz

Table 3: The experimental machine specifications. All of the CPU cores among the four computers can be fully employed by using the multi-threaded implementation, which significantly speeds up the evolution process.

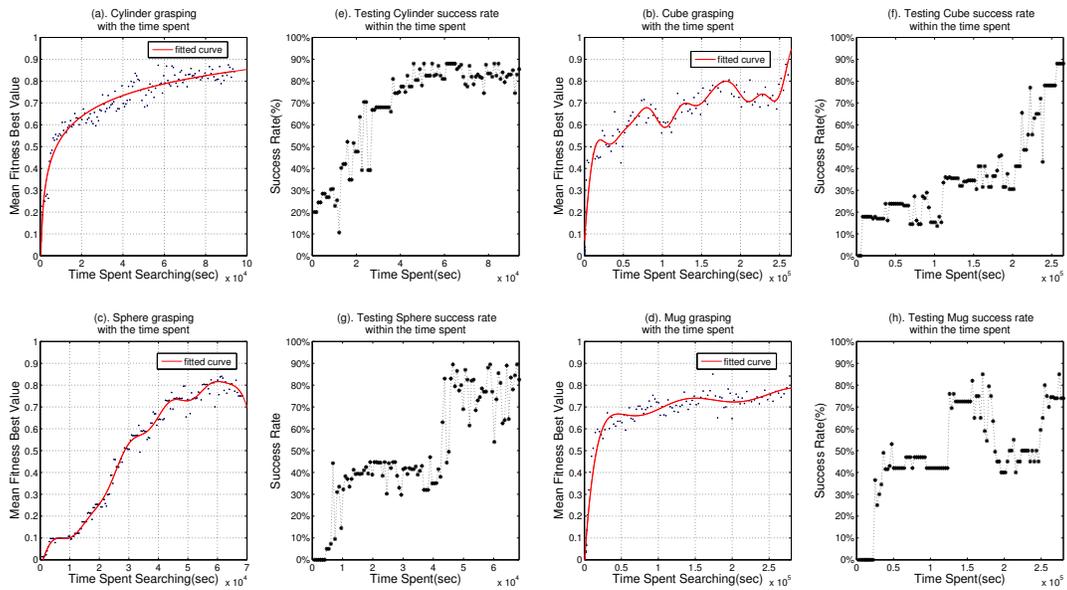


Figure 21: Training and testing performance with the different time spent searching. How fitness values increase over time is shown for each experiment. Plots (a) and (e) show a scenario with a single cylinder on a table, (b) and (f) a single cube on a table, (c) and (g) a single sphere on a table, (d) and (h) a single mug on a table. To evaluate whether Algorithm 1 benefits performance, (a)-(d) are training results, while (e)-(h) are testing outcomes. Fitness generally improves as training progresses. The conclusion is that although accuracy generally benefits from increased time, increased time may sometimes also there is risk overfitting to the training cases, as shown by the intermediate trough in test performance in (h).

to grasp a single cylinder, a single cube, a single sphere, and a single mug. Because the goal is to use only the best controller for the actual grasp, only overall best-case results are presented here. For each 10 minutes, the best-case value was recorded, and we repeated the same experiments three times, and the average values (the dots) and the fitting models are shown in Figure 21(a)-(d). These figures show how fitness values increase over evolutionary search. The underlying assumption is that larger fitness values implies better quality grasping.

To start the neuroevolution simulation, individuals in the population are initialized with random weights and a simple topology (i.e., in our case input nodes fully connected to four hidden nodes that are fully connected to the outputs), as shown in Figure 6. Because randomly generated policies generally do not cause the robot hand to approach the target objects, low fitness scores are expected. As time passes, Algorithm 1 leads the *Mekahand* to the right position/orientation toward the object. Accordingly, Figure 21(a), a cylinder, shows that initially the fitness value is low, and then after 80,000 secs, it reaches 0.8. Similar results appear in the other three experiments in Figure 21(b)-(d) (the cube, the sphere, and the mug). However, in Figure 21(b), (d), (the cube and the mug) the ANNs took 170,000 secs to get 0.8. In Figure 21(c), the sphere, the ANNs only used 50,000 secs to achieve 0.8. Two interesting results are noteworthy here. First, in the case of the sphere (Figure 21(c)), fitness dropped to 0.7 after 60,000 secs of training; one explanation is that sometimes overfitting occurs, which reduces accuracy. Second, in the case of the mug (Figure 21(d)), the fitness value plateaus after 100,000 secs. Such plateaus in evolutionary search often reflect a local optimum; the mug is likely difficult to distinguish from the other objects. Comparing the four figures, it can be seen that the neural networks trained on the ‘simple’ objects (Figure 21(a), (c), the cylinder and the sphere) were much easier trained than networks trained on objects of more

complicated geometry (Figure 21(b), (d), the cube and the mug). Figure 22 shows an example of an ANN evolved for 100,000 secs to grasp the cylinder.

The second set of experiments applies the NEAT algorithm on the same four target objects but in novel situations unseen in training, as discussed in Section 6.5. Here, the one-hundred best-trained neural networks that result from different search times were picked and tested with the same input. Each ANN was tested 200 times. The success rates shown in Figures 21(e)-(h) compare generalization of neural networks versus time spent in training. The first experiment in Figure 21(e), the cylinder, shows that as the search time increases, the success rate improves and after 45,000 secs, the success rate reaches and stabilizes at around 80%. Similar results are seen in Figure 21(f), (g) (the cube and the sphere). In Figure 21(f), (g), the maximum success rate is 88%. However, in Figure 21(g), the sphere, the success rate continues to oscillate after 40,000, even after the training fitness value has stabilized around 0.8. This may reflect the difficulty in finding a robust and general grasp for an entirely-curved surface; maybe more training time is required. Figure 21(h), the mug, illustrates that sometimes grasps for complex objects can be consistent, although there is some oscillation in test performance later in evolution.

7.3 Grasp Quality vs. Training Effort vs. Task Completion Time Tradeoff Evaluation

This set of experiments measures the success rate of grasps across a two-dimensional surface of tradeoffs, examining how success varies across both training effort and task completion time. The idea is to explore what minimum amount of training effort and task completion time is necessary to successfully grasp an object. The grasping procedures were implemented under the designed motion planner, whose sequence consists of five states and state transitions: (1) starting from initial to grasp coordinates, (2) grasp, (3) go to dropoff location, (4) place, and (5) go back to initial state. The different training stage of the evolved neural network were used to generate the grasping position and orientation of the *Mekahand* at step (1). A grasp was rated as successful if an object was grasped, lifted, and placed to another location from the initial position.

Figure 23 gives the complete grasping success rate with different task completion times and different training efforts. Each object was tested 10 times per task completion time and per training effort. For the implementation used, the grasping task completion time was measured ranging from 1 second to 10 seconds (the interval is 1 second) and the training effort was chosen ranging from 0 min to 1,200 mins (the interval is 120 mins). Therefore, the total number of object tests is $10 \times 10 \times 10$. Overall, these results indicate that when accumulated training effort is higher than around 600 mins, we could achieve 70% successful grasps rate. Interestingly, the task completion time has only a negligible impact on the success rate; even if only one second is allocated to task completion, the grasp quality remains nearly the same, indicating the controller has a stable design. The best match is depicted in Figure 23. If the success rate is set 70%, best performance results when the ANN is trained for 600 minutes, and the task completion time is 8 seconds. Figure 23 was meant to be used as a reference to indicate the expected success rate when pre-conditions are given, which is derived from the experiments, allowing decision-makers to make optimal tradeoffs between the grasping quality, the task completion time and the training effort in robotic control.

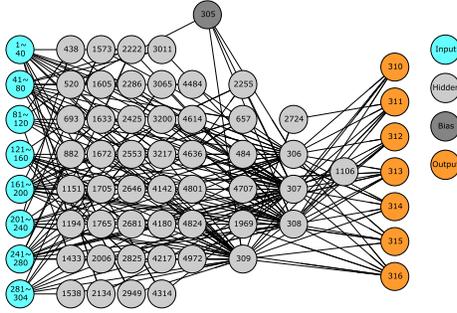


Figure 22: An example of a grasping network evolved by NEAT after 100,000 secs. The 304 Input nodes (20×15 pixels, plus a set of coordinate inputs, and object scale inputs) are located on the left, while the 7 output nodes are located at the right side. The light gray nodes in between are the evolved hidden nodes. The conclusion is that as time goes by, NEAT searches through increasingly complex networks to find one able to match the complexity of the grasping problem.

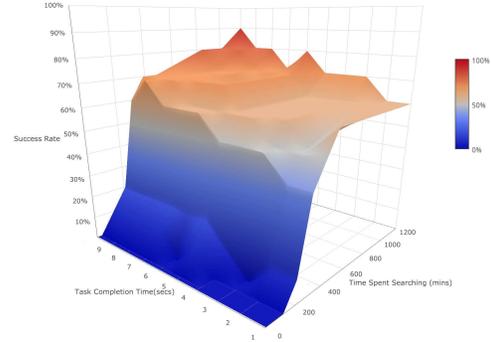


Figure 23: The grasp success rate vs. task completion time vs. training time. The x axis represents the task completion time, which ranges from 1 to 10 seconds, the y axis indicates the time spent searching ranging from 0 to 1,000 mins, and the z axis represents the resulting grasp success rate. In general, the results indicate that a 70% successful grasp rate could be achieved after around 600 minutes of training effort. Moreover, it helps highlight what task completion time is sufficient to successfully grasp an object given enough training effort.

8 LESSONS LEARNED AND DIRECTIONS FOR FUTURE WORK

There are several interesting observations we may infer from this tradeoff study. In the introduction, we asked whether the quality of grasping may exhibit a steep improvement after some critical number of training cases. As we can see from the results in Figure 21, the answer depends on the shape of the object being grasped. From the results, the fitness function does not seem to be a very reliable predictor of the success of a grasp, especially when training time is limited. A possible explanation of the results is that for objects of shapes that generate a search space with sharp local maxima, the NEAT algorithm does not perform well until the neural network evolved assumes a topology that is compatible with the geometry of the search space, i.e., the neural network implements the correct function for some choice of the edge weights. From that point onward, the search should converge quickly to the correct weight configuration. If this interpretation has merit, then it would make sense to focus on searching for the right topology rather than assigning the right weights in the early phase of learning. A good heuristic for selecting the right network topology for the robotic task of interest would be especially useful, perhaps by employing ideas from simulated annealing.

A more practical result of this paper is the observation that the imprecise computation framework is a good match for robotic skill acquisition. Inasmuch as the goal is to grasp an object well enough so that the object cannot be dropped easily by minor perturbation on the controller, it is important to delineate the boundary between a good enough grasp and one that may drop the grasped object. The tradeoff results are useful for this delineation. For future work, it is useful to extend the tradeoff analysis to more complex robotic tasks than grasping.

9 CONCLUSION

Future robotic systems will need to function in unknown and unstructured environments, under demanding timing constraints. This paper realizes a semi-autonomous remote robotic system that integrates body-compliant control, neuroevolution and real-time constraints. An important contribution of the paper is a tradeoff study in the design space of a representative robotic task, specifically, grasping of unknown objects in unstructured environments through imprecise computation. Such a design space typically includes three key dimensions: (1) the amount of training effort to teach the robot to perform the task, (2) the time available to complete the task once a command is given to perform it, and (3) the quality of the results of completing the task. The tradeoff results in this study indicate that training effort is a critical factor for attaining high-quality grasping. The trained model, after training, can then be reused without further modification. The work here can be viewed as offering a realistic basis for the scheduling work done by the real-time systems community in the past two decades. Importantly, beyond simulated tests, the results were also validated with a real robot (*Dreamer*) and contribute to the development of a systematic approach for designing robotic systems that can function in the challenging environments of the future, such as flexible manufacturing factories.

References

- [1] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *International Journal of Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- [2] Josh C. Bongard. The utility of evolving simulated robot morphology increases with task complexity for object manipulation. In *Artificial Life*, volume 16, pages 201–223, 2010.
- [3] Martin Buss, Hideki Hashimoto, and John B Moore. Dextrous hand grasping force optimization. *IEEE Transactions on Robotics and Automation*, 12(3):406–418, 1996.
- [4] Columbia University Robotics Groups. GraspIt! Simulator. <http://graspit-simulator.github.io/#>.
- [5] Mark 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, 1989.
- [6] Carlo Ferrari and John Canny. Planning optimal grasps. *Proceedings of the IEEE International Conference on Robotics and Automation*, (3):2290–2295, 1992.
- [7] Song Han, Aloysius K. Mok, Jianyong Meng, Yi-Hung Wei, Pei-Chi Huang, Quan Leng, Xiuming Zhu, Luis Sentis, Kwan Suk Kim, and Risto Miikkulainen. Architecture of a cyberphysical avatar. In *Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems*, pages 189–198, 2013.
- [8] Simon S. Haykin. *Neural Networks and Learning Machines*. Prentice Hall, Upper Saddle River, New Jersey, third edition, 2009.
- [9] Kaijen Hsiao, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Grasping POMDPs. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4685–4692, April 2007.

- [10] Kaijen Hsiao and Tomás Lozano-Pérez. Imitation learning of whole-body grasps. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5657–5662, October 2006.
- [11] Pei-Chi Huang, Joel Lehman, Aloysius K Mok, Risto Miikkulainen, and Luis Sentis. Grasping novel objects with a dexterous robotic hand through neuroevolution. In *2014 IEEE Symposium on Computational Intelligence in Control and Automation (CICA)*, pages 1–8. IEEE, 2014.
- [12] Pei-Chi Huang, Luis Sentis, Joel Lehman, Chien-Liang Fok, Aloysius K Mok, and Risto Miikkulainen. Tradeoffs in real-time robotic task design with neuroevolution learning for imprecise computation. In *Real-Time Systems Symposium, 2015 IEEE*, pages 206–215. IEEE, 2015.
- [13] Joshua M Inouye and Francisco J Valero-Cuevas. Anthropomorphic tendon-driven robotic hands can exceed human grasping capabilities following optimization. *The International Journal of Robotics Research*, 33(5):694–705, 2014.
- [14] Nick Jakobi. *Minimal simulations for evolutionary Robotics*. PhD thesis, University of Sussex, School of Cognitive and Computing Sciences, 1998.
- [15] Jeffrey Kerr and Bernard Roth. Analysis of multifingered hands. *The International Journal of Robotics Research*, 4(4):3–17, 1986.
- [16] David Kirkpatrick, Bhubaneswar Mishra, and Chee-Keng Yap. Quantitative steinitz’s theorems with applications to multifingered grasping. *Discrete & Computational Geometry*, 7(3):295–318, 1992.
- [17] Jens Kober, J. Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *International Journal of Robotics Research*, 32(11):1238–1274, July 2013.
- [18] Nate Kohl, Kenneth Stanley, Risto Miikkulainen, and Rini Sherony. Evolving a real-world vehicle warning system. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1681–1688, 2006.
- [19] Quoc V. Le, David Kamm, Arda F. Kara, and Andrew Y. Ng. Learning to grasp objects with multiple contact points. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 5062–5069, May 2010.
- [20] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [21] Joel Lehman and Risto Miikkulainen. Overcoming deception in evolution of cognitive behaviors. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 185–192. ACM, 2014.
- [22] Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for natural evolution novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.
- [23] Hod Lipson, Josh C. Bongard, Victor Zykov, and Evan Malone. Evolutionary robotics for legged machines: From simulation to physical reality. In *Proceedings of the 9th International Conference on Intelligent Autonomous Systems.*, pages 11–18, 2006.

- [24] Guanfeng Liu, Jijie Xu, and Zexiang Li. On geometric algorithms for real-time grasping force optimization. *IEEE Transactions on Control Systems Technology*, 12(6):843–859, 2004.
- [25] Jane W.-S. Liu, Kwei-Jay Lin, C. L. Liu, and C. W. Gear. Imprecise computations. *Mission Critical Operating Systems*, IOS Press, pages 159–169, 1992.
- [26] Jane W.-S. Liu, Wei-Kuan Shih, Kwei-Jay Lin, Riccardo Bettati, and Jen-Yao Chung. Imprecise computations. *IEEE*, 82(1):83–94, January 1994.
- [27] Yun-Hui Liu. Qualitative test and force optimization of 3-D frictional form-closure grasps using linear programming. *IEEE Transactions on Robotics and Automation*, 15(1):163–173, 1999.
- [28] Andrew P Martin. Increasing genomic complexity by gene duplication and the origin of vertebrates. *The American Naturalist*, 154(2):111–128, 1999.
- [29] Andrew Miller, Peter 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.
- [30] Andrew T. Miller and Peter K. Allen. Examples of 3D grasp quality computations. *Proceedings of the IEEE International Conference on Robotics and Automation*, 8:1240–1246, 1999.
- [31] Andrew T. Miller and Peter K. Allen. GraspIt!: A versatile simulator for robotic grasping. *IEEE Robotics and Automation Magazine*, 11(4):110–122, December 2004.
- [32] Andrew T. Miller, Steffen Knoop, Henrik I. Christensen, and Peter K. Allen. Automatic grasp planning using shape primitives. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1824–1829, September 2003.
- [33] Bhubaneswar Mishra, Jacob T Schwartz, and Micha Sharir. On the existence and synthesis of multifinger positive grips. *Algorithmica*, 2(1):541–558, 1987.
- [34] J. Molina-Vilaplana, J. Feliu-Batlle, and J. López-Coronado. A modular neural network architecture for step-wise learning of grasping tasks. *International Journal of Neural Networks*, 20(5):631–645, 2007.
- [35] Saigopal Nelaturi, Arvind Rangarajan, Christian Fritz, and Tolga Kurtoglu. Automated fixture configuration for rapid manufacturing planning. *Computer-Aided Design*, 46:160–169, 2014.
- [36] Leonardo M. Pedro, Valdinei L. Belini, and Glauco A.P. Caurin. Learning how to grasp based on neural network retraining. *Advanced Robotics*, 27(10):785–797, 2013.
- [37] Raphael Pelossof, Andrew T. Miller, Peter K. Allen, and Tony Jebara. An SVM learning approach to robotic grasping. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3512–3518, April 2004.
- [38] Deepak Rao, Quoc V. Le, Thanathorn Phoka, Morgan Quigley, Attawith Sudsang, and Andrew Y. Ng. Grasping novel objects with depth segmentation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2578–2585, 2010.

- [39] Nasser Rezzoug and Philippe Gorce. *Robotic Grasping: A Generic Neural Network Architecture*. Mobile Robots Towards New Applications, Berlin, Germany: Pro Literatur Verlag Robert MayerScholz, 2006.
- [40] Sebastian Risi, Charles E Hughes, and Kenneth O Stanley. Evolving plastic neural networks with novelty search. *Adaptive Behavior*, 18(6):470–491, 2010.
- [41] Máximo A Roa and Raúl Suárez. Grasp quality measures: review and performance. *Autonomous robots*, 38(1):65–88, 2015.
- [42] A. Saxena. *Monocular depth perception and robotic grasping of novel objects*. PhD thesis, Stanford University, 2009.
- [43] Ashutosh Saxena, Justin Driemeyer, and Andrew Y. Ng. Robotic grasping of novel objects using vision. *International Journal of Robotics Research*, 27(2):157–173, 2008.
- [44] Ashutosh Saxena, Lawson L.S. Wong, and Andrew Y. Ng. Learning grasp strategies with partial shape information. In *Proceedings of the 23th AAAI Conference on Artificial Intelligence*, pages 1491–1494, 2008.
- [45] L. Sentis, J. Park, and O. Khatib. Compliant control of multi-contact and center of mass behaviors in humanoid robots. *IEEE Transactions on Robotics*, 26(3):483–501, June 2010.
- [46] L. Sentis, J.G. Petersen, and R. Philippsen. Implementation and stability analysis of prioritized whole-body compliant controllers on a wheeled humanoid robot in uneven terrains. 35(4):301–319, November 2013.
- [47] Karun B Shimoga. Robot grasp synthesis algorithms: A survey. *The International Journal of Robotics Research*, 15(3):230–266, 1996.
- [48] T. Shin-ichi and M. Satoshi. Living and working with robots. *Nipponia*, (13):7–12, 2000.
- [49] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [50] Avishai Sintov and Amir Shapiro. An analysis of grasp quality measures for the application of sheet metal parts grasping. *Autonomous Robots*, 41(1):145–161, 2017.
- [51] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.
- [52] Kenneth O. Stanley and Risto Miikkulainen. Competitive coevolution through evolutionary complexification. *International Journal of Artificial Intelligence Research*, 21:63–100, 2004.
- [53] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [54] Marek Teichmann and Bud Mishra. The power of friction: Quantifying the “goodness” of frictional grasps. *Algorithms for Robotic Motion and Manipulation*, pages 311–320, 1997.

- [55] Jianwei Zhang and Bernd Rössler. Self-valuing learning and generalization with application in visually guided grasping of complex objects. *Robotics and Autonomous systems*, 47(2):117–127, 2004.