Physically-based Grasp Quality Evaluation under Pose Uncertainty

Junggon Kim, Kunihiro Iwamoto, James J. Kuffner, Yasuhiro Ota and Nancy S. Pollard

***Abstract*—Although there has been great progress in robot grasp planning, automatically generated grasp sets using a quality metric are not as robust as human generated grasp sets when applied to real problems. Most previous research on grasp quality metrics has focused on measuring the quality of established grasp contacts after grasping, but it is difficult to reproduce the same planned final grasp configuration with a real robot hand, which makes the quality evaluation less useful in practice. In this study we focus more on the grasping process which usually involves changes in contact and object location, and explore the efficacy of using dynamic simulation in estimating the likely success or failure of a grasp in the real environment. Among many factors that can possibly affect the result of grasping, we particularly investigated the effect of considering object dynamics and pose uncertainty on the performance in estimating the actual grasp success rates measured from experiments. We observed that considering both dynamics and uncertainty improved the performance significantly and, when applied to automatic grasp set generation, this method generated more stable and natural grasp sets compared to a commonly used method based on kinematic simulation and force-closure analysis.**

***Index Terms*—Grasp quality evaluation, object dynamics, pose uncertainty**

I. Introduction

Grasping objects with robot hands reliably and stably is one of the key goals of robust manipulation, but it is still challenging to achieve it in a real environment. One way to achieve more robust grasping is to evaluate the grasps in simulation. This approach has been extensively studied over the past decades.

Most previous research on grasp quality evaluation focused mainly on measuring the quality of an established final grasp configuration, or the situation where the robot hand is already holding an object with contacts. In practice it is difficult to reproduce the same final grasp configuration with a real robot hand due to limitations in sensing and control of the robot system. In most cases, the object may move unexpectedly due to finger contacts during grasping. This could cause a catastrophic failure such as dropping, or result in a different robot hand configuration and different contacts from the originally planned grasp, making the quality evaluation less informative.

Different from most previous work, we focus more on the grasping process than the final grasp configuration. We want to estimate the likely success or failure of a grasp in the real environment using a simulation technique so that we can use more robust grasps and avoid grasps that are likely to fail in practice. Grasping involves rapid change in contact between the hand and the object, and this also changes the position and orientation of the object being grabbed. Thus, it is natural to try using a dynamic simulation technique, instead of kinematic simulation that has been commonly used in many grasp quality measure literature, to predict the result of grasping more correctly.

Most existing solutions to dynamic simulation have put focus on improving the plausibility of the resulting motions or speeding up the computation time by developing better techniques for modeling and handling the equations of motion. However, efficacy of using dynamic simulation in evaluating the quality of a grasp has been rarely studied in the grasp community, even though the grasping process cannot be well described with a static analysis model. In fact, it is still an open question how much we can trust the result of dynamic simulation especially when the simulation involves complex interactions through rapidly changing contacts because there are too many uncertainties in the real world that cannot be exactly modeled within the simulation.

In this article we explore the question of "*Can we use dynamic simulation to estimate the probability of success or failure of a grasp in the real environment?*" which is a fundamental question in both the grasp and simulation communities, and what we found is that using dynamic simulation improves the performance of grasp evaluation in predicting actual grasp success rates if uncertainty is considered within the simulation. This is our main contribution.

We define a grasp as the combination of a relative pose (position and orientation) of the hand to the object and the finger joint angles prior to grasping. A grasp for an object and a robot hand is regarded as successful when the hand can grab the object securely by closing the fingers from the grasp, and as a failure if the robot drops or loses the object during the grasping process. Note that the grasp definition actually represents a 'pre-grasp' prior to closing the fingers, and this form of definition is suitable to a data-driven grasp planning approach where the robot chooses a grasp having a feasible trajectory from a precomputed set of good grasps, moves the robot hand to a particular place specified by the grasp, and finally closes the fingers to grab the object.

Since most of the existing grasp quality measures compute the quality score based on contact information, many grasp analysis tools such as GraspIt! [24] and OpenRAVE [10] run a simple kinematic grasping simulation to obtain relevant

Junggon Kim and Nancy Pollard are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA {junggon,nsp}@cs.cmu.edu

Kunihiro Iwamoto and Yasuhiro Ota are with the Partner Robot and Advanced Engineering Group, Toyota Motor Engineering and Manufacturing North America, CA, USA. {kunihiro.iwamoto,yasuhiro.ota}@tema.toyota.com

James Kuffner is with Google Inc., CA, USA. kuffner@google.com

contact points to evaluate a grasp. They close the fingers until touching the object or reaching the joint limits while the object remains at the same location even after the collision by the fingers. However, such an existing method, i.e., the combination of kinematic grasping and an existing grasp quality metric, performs poorly in predicting the success or failure of actual robot grasping. In many cases the method misjudges a good grasp as a bad one, or vice versa. For example, the two grasps shown in Figure 1 for a mug and a watering can are two of the best grasps chosen by the existing method, but they look fragile and would not perform well in the real environment. On the contrary, the grasp shown in the right of the figure was discarded because it did not form force-closure under the static object assumption though, in the real environment, it would have resulted in a stable power grasp after the object's location is changed by the contacts.

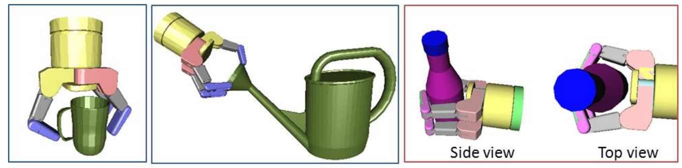

Fig. 1. **(Left, middle)** Two high ranking grasps chosen by a typical existing method consisting of a kinematic grasping simulation with static object assumption and a force-closure based quality metric [11]. However, they look fragile and would not work well in the real world. **(Right)** A grasp discarded by the existing method. Due to the static object assumption, the object remains in its original location even after the finger contacts. As a result, the palm still does not support the object and the grasp does not form a force-closure. In the actual grasping, however, this would have resulted in a stable power grasp.

In this paper we discuss how to improve upon this grasp quality evaluation method using dynamic simulation. There are many factors that can affect the success or failure of a grasp, such as dynamics of the system including the object and the hand mechanism, uncertainty in the modeling and sensing, and the control algorithm for finger closing. However, taking all of the factors into account in the simulation is difficult to implement and to validate. In this study, we focus on investigating the effectiveness of considering dynamics in grasp evaluation rather than building a dynamic simulator for accurately reproducing grasping. Thus, our approach is to concentrate on a small number of key factors directly affecting grasp success or failure, and to remove other factors from consideration.

More specifically, we investigate the effect of considering two missing factors in the kinematic simulation – object movement during grasping and pose uncertainty – on the performance in estimating the actual grasp success rates measured from experiments. Though the two factors always exist in the real world and can affect the grasping result significantly, they have been rarely considered at the same time in most of previous research on grasp quality evaluation. In our simulation, full 3D dynamics of the object, which is assumed as a rigid body, is considered to capture its motion interacting with the hand during grasping. The uncertainty in the object pose is handled by running the dynamic simulation multiple times where each simulation starts from a slightly different initial condition sampled from an error model. To simplify the problem, the dynamics of the robot hand mechanism is ignored. A simple penalty-based contact model is used to consider the interaction between the object and the hand during grasping. We also assume a simple open-loop grasp controller where the closing speed is adjusted based on grasp contact forces. Better estimation of the success rate of such an open-loop grasp is useful in itself for systems that rely on a relatively simple finger closing mechanism for grasping. Our approach, however, may also be used in conjunction with more sophisticated sensor feedback driven grasping algorithms to remove poorly performing grasps from consideration.

In our prior work [19], we evaluated the quality of grasps with a new method considering the two key factors, and showed its effectiveness in estimating actual grasp success or failure through an experiment with simple objects. We extend the work by refining the suggested measuring algorithm, adding more algorithms to the comparison, and testing with more complicated objects. We also apply the method to automatic grasp set generation and discuss the benefits of using it over the existing method.

The rest of this article is organized as follows: We briefly review relevant previous work in Section II, and introduce our simulation based methods for grasp evaluation in Section III. The robot experiment for measuring the actual grasp success rates for selected grasps and objects is described in Section IV. We compare the simulated grasp quality score data with the measured success rates, and analyze the performance of the tested evaluation methods in predicting the actual grasp success rates in Section V. We apply the simulation-based method to grasp set generation and discuss the benefit of using the new method over existing methods in Section VI, and conclude this work in Section VII.

II. Previous Work

Much previous work on grasp quality metrics has focused on analyzing the 6-dimensional space spanned by contact wrenches. Li and Sastry [22] suggested using the smallest singular value of the grasp matrix, which relates the fingertip forces and the net wrench applied to the object, and the volume of the wrench space as quality metrics. They also proposed a task oriented quality measure to consider the type of task to be done with the grasp. Ferrari and Canny [11] suggested using the largest disturbance wrench that can be resisted in all directions by the contacts. Geometrically, this corresponds to the radius of the largest ball, centered at the origin, which is contained within the convex hull of the unit contact wrenches. This force-closure based metric is one of the most popular methods for measuring grasp quality and has been implemented in many systems for grasp analysis such as GraspIt! [24]. Ciocarlie and Allen [8] proposed a metric considering the distances between predefined contact points on the hand and the object to assess the quality of a pre-grasp which is very close to, but not in contact with the object, and applied this to online dexterous grasp planning using optimization on a low-dimensional subspace. Balasubramanian et al. [2] investigated grasp quality measures that may be derived from human-guided robot grasps and reported that the wrist orientation for the highly successful human generated grasps tends to be aligned to the principal axes of the object much more closely than for the GraspIt! grasps. Metrics considering certain geometric relations of the contact points have also been proposed in [5], [25], [27]. We refer readers to [32] for a nice summary of a variety of grasp quality metrics.

Grasp quality measures have many applications such as finding an optimal grasp [11], [18], [22], [25], and generating grasp sets [10], [15]. In this paper, we primarily consider the application of grasp set generation, although our findings are relevant to other applications as well. A large grasp database containing grasps of various objects was built using GraspIt! [15]. Other manipulation planning tools such as OpenRAVE [10] also provide a function to sample grasp candidates, test them with a metric, and build a grasp set automatically in order to use the precomputed grasp set in motion planning algorithms such as RRT-Connect [20]. Most of the existing methods for automatic generation of a grasp set obtain the contact points of the grasp using a simple kinematic simulation of finger closing with a static object assumption, i.e., the object remains at the same position even after collision with the fingers. However, the assumption does not hold in many practical situations because the object can move significantly in response to the collision during grasping and in such cases the existing metrics may not give us useful information on grasp quality. To handle the issue, we need to consider dynamics of the object in the simulation for evaluating grasp quality.

There are many solutions to obtain physically plausible interactions between a robot hand and an object [1], [10], [21], [23], [30], but most of them focused on building simulation tools and did not investigate the effectiveness of using dynamic simulation in evaluating grasp quality. Goldfeder et al. [13] employed an approximate dynamic simulation, in addition to the static simulation, where object response due to forces applied by the hand is captured without considering other environment forces such as the supporting surface, but they did not discuss the efficacy of considering dynamics in their grasp quality analysis in detail. Zhang et al. [35] explored whether a dynamic simulation model can predict the actual grasping process accurately, using a planar grasping testbed equipped with a linear actuator as a thumb and three fixed pins as fingers. They showed that simulation with a well-calibrated model can be effectively used as a surrogate for real experiments in their 1-dof planar grasping setting. In this article we explore the efficacy of using dynamic simulation in predicting the likely success or failure of a grasp in 3D environment.

Another issue in evaluating grasp quality is data uncertainty. There is an extensive body of research on action planning in the presence of significant uncertainty such as the work by Brost [4], and Goldberg and Mason [12]. However, we assume that the robot is capable of estimating object pose with moderate error. Zhen and Quian [36] investigated how small uncertainty in the friction coefficients and contact locations affects grasp quality. Christopoulos and Schrater [6] similarly incorporated shape uncertainty into grasp stability analysis of two-dimensional planar objects by considering the effect of small changes in contact force position and direction. Goldfeder et al. [13] handle shape uncertainty by cross testing grasps with alternative shapes that are nearest neighbors to a given model. Hsiao et al. [16] introduced a method considering uncertainty in object shape and pose data by combining the data from a set of object detection algorithms using a probabilistic framework to find an optimal grasp.

Running multiple simulations with sampling has also been used to consider pose uncertainty in evaluating the quality of a grasp. OpenRAVE [10] computes grasp repeatability statistically by iterating kinematic grasping with randomly sampled object pose deviations and identifies the grasp as fragile if the deviation of the gripper's surface points is significantly larger than the deviation of the object's surface points. Most recently, Weisz and Allen [33] incorporated pose uncertainty into the static grasp quality analysis by computing the probability of force-closure in the presence of pose error. They sampled pose error uniformly from a 3-dimensional error model representing an object on a support plane and applied the existing force-closure analysis to compute the probability. In our prior work [19], we took a similar sampling approach to consider pose uncertainty. However, we used a probabilistic distribution model for the pose error and incorporated object dynamics along with different metrics into the grasp quality evaluation.

Notably, the effects of uncertainty and dynamic effects together have not been examined in detail, with simulation results compared to experiments in prior research. In [19] we show that considering both effects together is critical. In the present manuscript we build upon and elaborate those results.

III. Grasp Quality Evaluation

In this section we describe the elements we used to evaluate the quality of a grasp in detail. Open-loop grasping is simulated using our in-house physically-based grasp simulator and the quality of the grasp is evaluated based on the simulation result with a few measures which will be detailed below. In order to consider the object pose uncertainty, the simulation is repeated multiple times for each grasp where, at each trial, the simulation starts from a slightly different initial condition, and then the evaluated scores are averaged.

A. Grasping Simulation

In our simulation, after placing an object on a planar surface and moving the robot hand to a particular place specified by the grasp definition, we close the fingers to grab the object with closing speed adjusted based on the magnitude of the calculated contact forces. After all fingers have been closed, the hand is lifted a certain distance in order to see if the grasp can hold the object without ground support.

Full 3D rigid body dynamics of the object is considered to capture its motion interacting with the robot hand and the planar surface during the grasping process. More specifically, at every time step, the acceleration of the object is calculated from the current state of the object and the contact forces, and then is integrated to obtain the state of the object at the next time step. The frictional contact forces between the object and the hand, and between the object and the ground, are computed using a penalty-based method by Yamane and Nakamura [34]. The hand geometry is modeled as a set of uniformly distributed points, and a point-triangle collision detection algorithm is used to find the contact points.

We use a kinematic hand model consisting of the rigid links and the finger joints. The joints are driven by motors to close the fingers, and the closing speed is adjusted depending on the magnitude of the motor torque by

$$v = \begin{cases} v_0(1 - \frac{\tau}{\tau_m}) & (\tau < \tau_m) \\ 0 & (\tau \geq \tau_m) \end{cases} \tag{1}$$

where the motor torque τ is obtained by converting the contact forces using Jacobian matrices and the nominal closing speed v_0 and the torque limit τ_m are set by users. As the motor torque increases due to the forces at the contacts, the motor speed is reduced linearly. Note that the motor torque is highly approximated by ignoring the dynamics of the mechanical hand and the actual control system. The method, however, is still effective in simulating a simple open-loop closing mechanism, that is often found in many robotic systems, while saving the large amount of effort that would otherwise be required to obtain an accurate model of the complicated hand dynamics.

If the motors continue to close the fingers after touching the object, the fingers penetrate the object a little bit more at the next time step. This increases the penalty-based contact forces, and accordingly, decreases the closing speed. This mechanism eventually stops the closing at some point when the motor torque exceeds a given limit. Optionally, we can set a break-away torque so that the finger can keep closing with its outer joint only while the inner joint remains in place when the inner link has been blocked by the object. Finally, if all fingers have been closed or do not move for more than some period of time, the hand starts lifting the object to a certain point along a predefined trapezoidal velocity profile. During lift-up, the finger closing mechanism is still running, so the fingers keep squeezing the object and can close further when the object moves due to the changing circumstances. In our simulation the fingers are not back-drivable.

We monitor the simulation results such as the object pose and the contacts at each time step. After the simulation is done, we evaluate the grasp quality based on the gathered simulation data using the measures described below.

B. Grasp Quality Measures

As mentioned before, our focus in this study is on predicting the likely success or failure of a grasp using simulation. A grasp is regarded as successful when the hand can grab the object securely by closing the fingers from the grasp, and as a failure if the robot drops the object during the grasping process. Thus, one obvious way to evaluate the success or failure of a simulated grasp is to check if the hand still holds the object after the grasping simulation by counting the contacts between the object and the hand (Measure B). However, we are also interested in developing a new measure based on the monitored simulation data such as pose deviation during grasping (Measure C) and investigating if such a measure is effective in estimating grasp success or failure. Finally, a popular existing measure is employed in our analysis for comparison purposes (Measure A).

- **Measure A:** Maximum disturbance wrench
  Computing the maximum disturbance wrench that can be resisted by the contacts (Ferrari and Canny [11]) is one of the most popular ways to evaluate the quality of a grasp, and this metric has been used in many grasp analysis tools [10], [24]. Once a final grasp configuration with contacts has been obtained after the grasping simulation, we compute the minimum distance to the boundary of the convex hull of the unit contact wrenches to obtain the maximum disturbance wrench that can be resisted. We use OpenRAVE [10] to perform the computation. Note that even with this metric, our approach differs from the approaches used in the existing tools because we consider the dynamic motion of the object, which may be substantial, during grasping.

- **Measure B:** Number of contact links
  One simple way to evaluate the success or failure of a simulated grasp is to check if the hand still holds the object by counting the contacts between the hand and the object after the grasping process. We judge that a simulated grasping has failed if the object was out of the hand or it had contacts with less than two hand links at the end of grasping. If the object was held within the hand and supported by three or more links, we regard the grasp to be successful. If the object was supported by only two links so that it may dangle, we give a half credit to the grasp. Note that we count the number of contact links, not the number of contact fingers, so grabbing with only two fingers can still get the full credit in some situations. After the lift-up process, our system counts the number of contact links, and measures the quality score using the following 3 step scoring system:
  - 1: The hand is holding the object with three or more contact links.
  - 0.5: The hand is holding the object with only two contact links, so the object is likely to dangle.[1]
  - 0: The hand failed in grasping the object.

  The metric is easy to implement in the simulation system and also convenient to apply to the experiment for measuring the actual grasp success rates (Section IV). The 3-step scoring system does not return a continuous score so it might not discriminate a good grasp from bad ones very well or vice versa. We make up for this by averaging the score values obtained from multiple simulations (Section III-C). The score value could also be weighted by the number of contact points at each link or the contact properties such as the normal direction for possible improvement, but those were not tested in

[1] Although we did not consider this, in case of non-rigid object or fingers, two contacts can form a force-closure or stable grasp due to the non-zero contact area. Refer to [7] to see how to incorporate soft contact into a force-closure based grasp analysis.

this study.

- **Measure C:** Pose deviation
  Here we assume that a grasp is better if it makes the object move less during grasping because unintended object movement caused by finger contacts is likely to increase the chance of failure in grasping. At each time step, the relative pose of the object to the hand is monitored and the pose deviation from the initial pose is calculated. Then, a grasp quality score is obtained by:

$$q = \begin{cases} 1 - \frac{\delta}{L} & (\delta < L) \\ 0 & (\delta \geq L) \end{cases} \quad (2)$$

  where δ is the pose deviation and L is a deviation limit which is introduced to normalize the deviation and defined by user.
  We consider the deviation in the position and orientation separately and compute them with

$$\delta_p = \|p_{\mathrm{com}} - \bar{p}_{\mathrm{com}}\|, \quad \delta_R = \|\log(\bar{R}^T R)\| \quad (3)$$

  where $p_{\mathrm{com}} \in \Re^3$ and $R \in SO(3)$ denote the relative center of mass position and the orientation of the object with respect to the hand coordinate system, and the bar symbol represents the reference value for measuring the deviation. Note that δ_p and δ_R are invariant under change of coordinate frames for the hand and the object.
  We can apply the quality measure in many ways. One way, which is used in this paper, is to use the pose deviation at the end of grasping which measures the difference between the planned and final object poses with respect to the hand. In our analysis (Section V-A), we set the deviation limit L to 5 cm and 30 deg for position and orientation respectively[2], and the obtained quality scores for position and orientation were averaged. Another way is to use the maximum pose deviation during grasping, which is more conservative. We can also try using the pose deviation measured in a particular period of time, such as the lift-up process as tested in our prior work [19].

C. Object Pose Uncertainty

In the real environment, we do not know the exact positions and orientations of the hand and the object due to sensor uncertainty. Thus, the actual hand pose relative to the object before grasping is always different from the ideal pose specified by the grasp definition, and this error can affect the success or failure of the grasp significantly.

There are two major sources of sensing error in the robot system we used in our experiments – a vision-based object pose estimation system and a cable-driven robot arm manipulator[3]. The actual uncertainty is affected by many factors. For example, the object pose error from the vision-based system may vary depending on the position of the object. The arm manipulator is known for its high backdrivability, but the cable driven mechanism makes it difficult to achieve precise position control, so the end-effector pose error may vary depending on the joint configuration and loading and also in time even after the calibration is done. Thus, it is very difficult to precisely measure the amount of uncertainty while considering all the possibilities that can arise in the real environment.

One way to consider pose uncertainty is to run multiple grasping simulations starting from slightly different initial conditions for each grasp where the initial condition is set by sampling from an error model representing the pose uncertainty. OpenRAVE [10] computes grasp repeatability by iterating kinematic grasping with randomly sampled object pose deviations. Weisz and Allen [33] used a regular sampling from a bounded 3-dimensional parametrization space (x, y, θ) to obtain the probability of obtaining a force-closure grasp, but applying a regular sampling to the full 6-dimensional pose space error is computationally expensive. Note that existing approaches assume the object is stationary during the grasping process and do not consider the effect of the pose error on the movement of the object during grasping.

We use a Monte Carlo method. For each grasp, the grasping simulation is repeated multiple times where the initial condition is set by randomly sampling from a probabilistic pose error model, and the evaluated grasp quality scores are averaged. We assume the pose uncertainty follows a normal distribution in the 6-dimensional pose space. In most of the simulation results shown in this paper, we used $(e_p, e_R) = (5\mathrm{mm}, 5\mathrm{deg})$, for every object, where e_p and e_R denote the expectations of the half-normal distributions for the positional and rotational pose errors respectively.[4] We note that, however, we have not measured the actual pose uncertainty in our robot system. Our focus in this study is on investigating the qualitative effect of considering pose uncertainty rather than on making a physically correct model of the real world (see Section V-A). We will also discuss how the change of the parameters of our uncertainty model affects the simulation result in Section V-B.

One of the main issues in using a Monte Carlo method is to set an appropriate sample size. If the sample size is too small, the result will not reflect the underlying uncertainty correctly. If the size is chosen too large, the method can become computationally too expensive. In order to determine an appropriate sample size for our grasp quality evaluation, we observed how the estimated quality scores and their range of error change with iterations. Figure 2 shows the error bars of the evaluated grasp quality scores for three test grasps. For each grasp, we ran 20 sets of independent sampling processes, and each process iterated up to more than 1000 times. At each

[2] We set 1 cm and 10 deg for the deviation limits in our prior work [19]. In this work we increased the values to 5 cm and 30 deg because we added larger objects such as Long box and Watering can in the experiment (Table I). However, it is unclear to us whether using different values (e.g., by scaling to the size of the object) for each object would be better than using constant values as we did in this study.

[3] WAM Arm from Barrett Technology Inc.

[4] In our prior work [19], we used $(e_p, e_R) = (1\mathrm{cm}, 6\mathrm{deg})$ in the analysis where the parameter values were determined based on the reported accuracy of the vision-based object pose estimation system (MOPED-1V [9]) used in the experiment. However, this does not include other sources of sensing error such as manipulator calibration error which we did not measure. In this paper we examined nine different sets of the parameters $\{(5\mathrm{mm}, 5\mathrm{deg}), (5\mathrm{mm}, 10\mathrm{deg}), \cdots, (15\mathrm{mm}, 15\mathrm{deg})\}$ to see how the parameter setting affects our grasp quality analysis. See Section V-B for more detail.

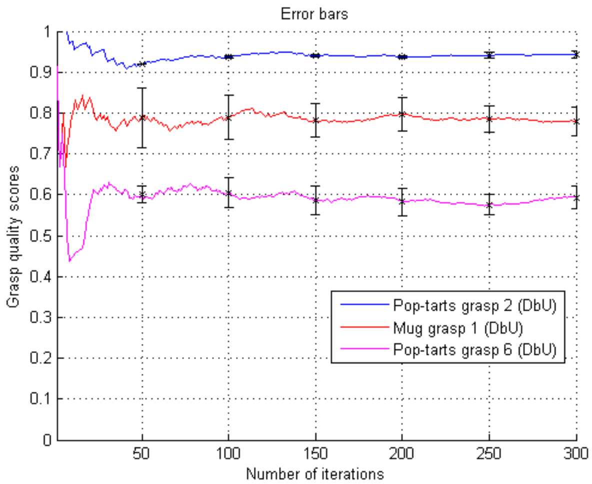

Fig. 2. The mean and standard deviation of the grasp quality score over the iterations.

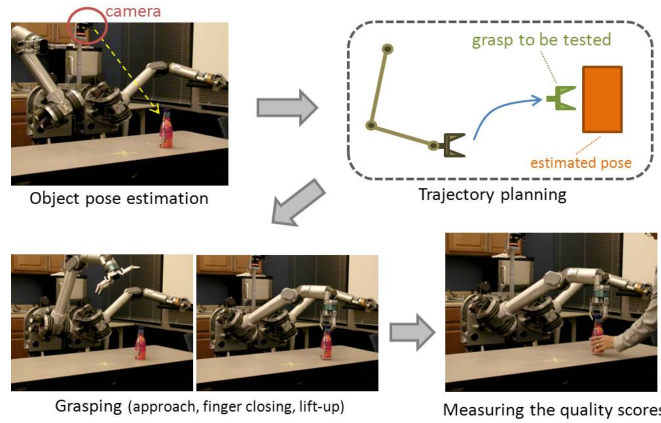

Fig. 3. Procedure of our experiment

iteration in a process, the grasp quality score was obtained by averaging the quality values in the current and all previous iterations. Then, the mean and standard deviation of the quality scores of all 20 processes were computed at each iteration, and marked in the figure at every 50 iterations. The grasp quality score appears to start getting stabilized from 50 iterations, and based on this observation, we set 100 as the sample size for testing a grasp under uncertainty. Note that, however, there is no specific rule of thumb for deciding the sample size. For example, considering a substantial reduction in the standard deviation between 100 and 150 samples in the middle graph of Figure 2, one can be more conservative by choosing a larger number (e.g., 150) as the sample size.

IV. Experiment

In order to evaluate the performance of the simulation-based method described in Section III in predicting the likely success or failure of a grasp, we measured actual grasp success rates experimentally. In our experiments we used an open-loop grasp with a simple finger closing mechanism. After the robot grabs an object with the grasp under testing, we manually inspect the final grasp configuration with our own scoring systems and repeat this multiple times to obtain an averaged score value which is regarded as the actual grasp success rate of the grasp.

We conducted the experiments with HERB, a service robot equipped with Barrett WAM arms and Barrett hands [31]. The procedure of our grasping experiment is shown in Figure 3. At each trial for testing a particular grasp, we place an object at a random location on a table. The robot estimates the pose (or the position and orientation) of the object using its own vision-based system [9]. Based on the estimated object pose, an RRT-based robot planner creates a trajectory to the grasp and executes the robot to approach the object along the trajectory. When the robot hand has reached the grasp pose (i.e., the pre-grasp), we command the robot to close the fingers to grab the object and then lift it up. Finally, we examine the grasp manually as described below and mark the score of the grasp. For each grasp of an object, we repeated this 10 times to obtain an averaged quality score of the grasp.

In the experiments we focused on the success or failure of grasping. If the robot hand has failed in grasping the object (e.g., dropping, which is quite obvious to judge by the human operator), we give the score of 0 to the grasp. If the hand successfully grabbed the object and the final grasp is secure, we give the score of 1 to the grasp. But, in reality, there exist situations that are hard to judge success or failure — in fact, even though the robot can grab an object, it is not always easy to say that the grasp is secure or fragile.

People are the best experts at grasping, and we believe that humans are also expert at judging how good or secure an established grasp would be in performing a task such as moving the object into another place. Perhaps, a person can make a best decision on this when he or she is allowed to touch and jiggle the object and the robot hand with his or her own hand interactively and feel the stability of the grasp, which is reflected in our interactive inspection method described below. We also tested a simplified method which counts the number of contacts visually without touching the object. The visual inspection is conceptually the same as Measure B described in Section III-B, and can be regarded as a possible surrogate for the interactive inspection as discussed later (Figure 5).

- **Visual inspection with a 3-step scoring system:** We use the same 3-step scoring system (Measure B) described in Section III-B. We regarded a final grasp with three or more contacts as a success and gave the score of 1 to the grasp. In case the fingers grabbed the object with only two contacts so that the object could possibly dangle, we gave 0.5 to the grasp. If the hand failed in grasping, we gave 0 to the grasp. We inspected the final grasp with our naked eyes without touching the object to count the number of contacts. Note that, again, we are counting the number of contact links, not the number of contact fingers.

- **Interactive inspection with a 5-step scoring system:** We deliberately touched and jiggled the object grabbed by the robot hand with human hand to feel the stability of the

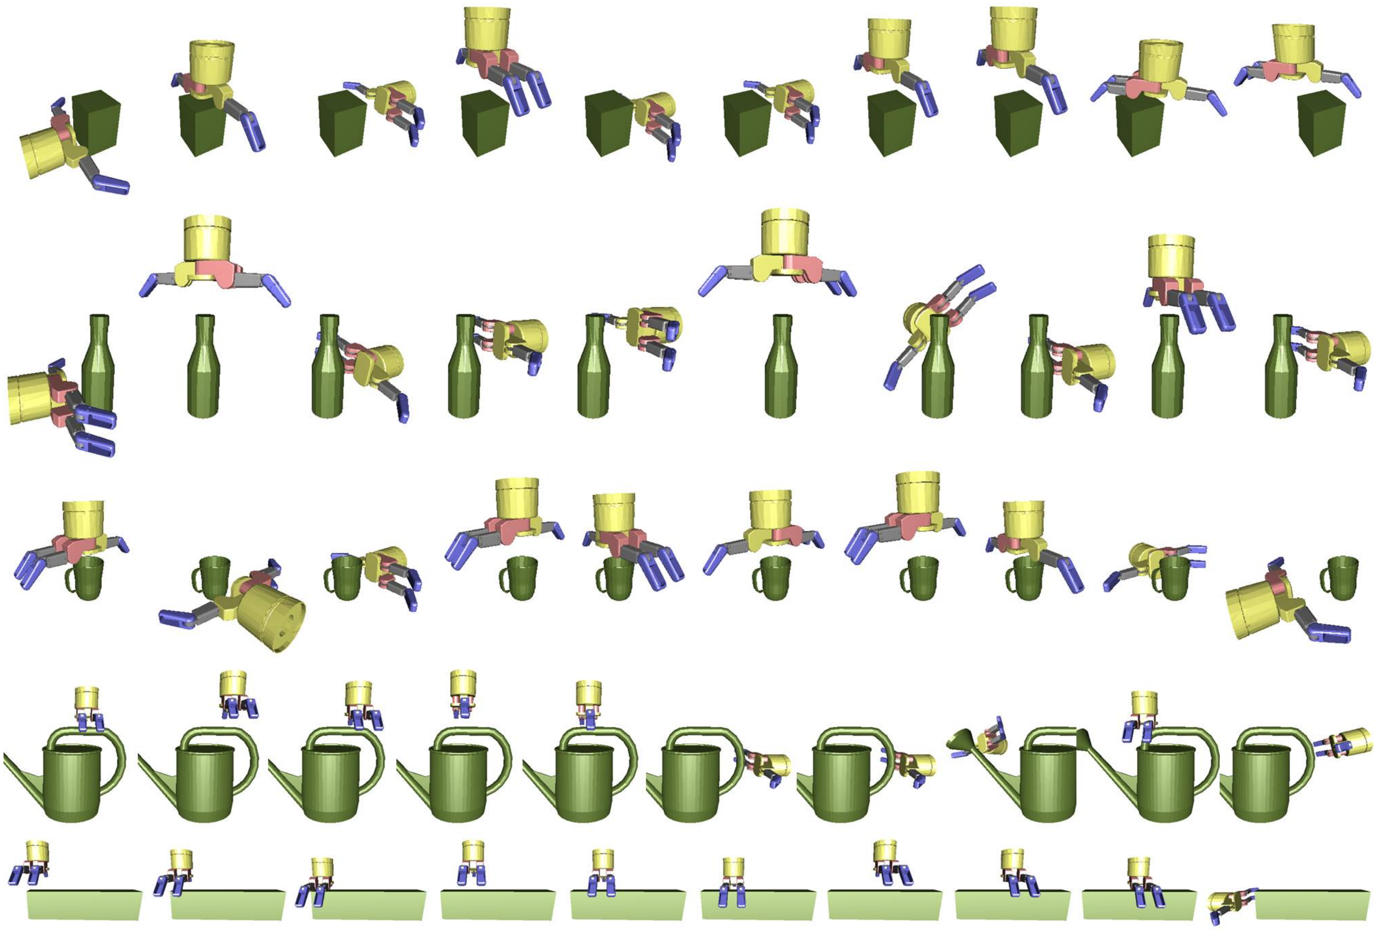

Fig. 4. The grasps used in the experiments. They were also used in our performance analysis described in Section V-A.

grasp, and marked the quality score using the following 5-step scoring system.

- 1: Stable grasp (unmovable by the small disturbance force).
- 0.75: The object moves by the disturbance force, but will not drop.
- 0.5: The object is movable and droppable by the disturbance force.
- 0.25: The grasp is fragile (won't be able to carry the object with the grasp).
- 0: Grasping failed.

At least two people participated in the measurement and achieved consensus at every trial to try to keep objectivity in scoring. The forces applied to the object by the human hand for inspection were about 2 N or less.

One advantage of using manual inspection methods is that we can directly use human intuition and experience in evaluating the quality of a grasp. Such a direct use of human intelligence in decision making often leads to a better and reliable result.[5] However, this requires a significant amount of human labor, and also, the scoring might not be objective though we tried to attain objectivity as much as possible. To

TABLE I
THE OBJECTS TESTED IN OUR EXPERIMENTS

Objects	Description	Mass (g)
Pop-tarts	A paper box filled with contents.	175
Fuze bottle	A plastic juice bottle filled with water.	595
Mug	A plastic mug.	80
Watering can	An empty plastic watering can.	430
Long box	A paper box filled with contents.	880

alleviate these issues, one can employ an automatic evaluation process such as the one used by Morales et al. [26] where they used three consecutive shaking movements of the hand to test the stability of a grasp in conjunction with tactile sensors for checking whether the object has been dropped by the shaking.

Five objects have been tested in the experiments (Table I). For each object, we generated a set of force-closure grasps using OpenRAVE and manually chose the 10 grasps, shown in Figure 4, having different approaching directions and distances from the object. The mug, long box, and watering can were

[5]For example, Balasubramanian et al. [2] showed human-guided grasps often perform better than the best grasps chosen by GraspIt!.

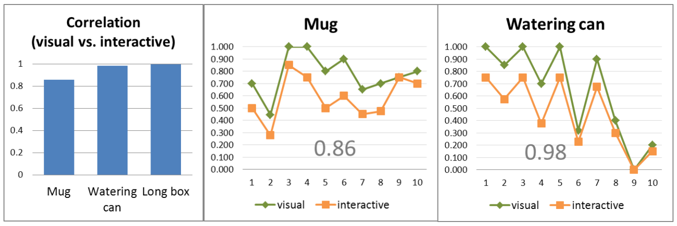

Fig. 5. **(Left)** Correlation between the two sets of experimental grasp quality scores. **(Middle, right)** The quality scores of the grasps for Mug and Watering can. The numbers in gray are the correlation coefficients.

tested with both 3-step and 5-step scoring systems while Pop-tarts and Fuze bottle were tested with the 3-step scoring system only. Figure 5 shows the two sets of the experimental grasp quality scores with their statistical correlation values.[6] Though their magnitudes are somewhat different, the two score sets show a very good consistency. For the big objects (long box and watering can), the correlation coefficients between the two score sets were very high while the coefficient got a little bit lower for the small and light object (mug). We used the experimental data obtained from the visual inspection (3-step scoring system) as the ground truth when we evaluate the performance of the simulation-based methods (Section V) because we used it for all objects. However, because of the good consistency between the two scoring systems, we also got a similar result when we replaced it with the interactive inspection data.

V. Results

In this section we investigate the performance of simulation-based methods in predicting the actual grasp success or failure. We calculate the quality scores using the measures described in Section III-B and obtain the correlation coefficients between the score sets from the simulations and the experiment. We also examine the effect of changing parameter settings for the uncertainty model and finger closing on the performance of the methods.

A. Performance Analysis

We first examine an existing method, or the combination of the kinematic grasping and a force-closure based measure, which is one of the most popular methods that has been used for generating grasp sets automatically. According to our study, however, the method turned out to be poor in predicting the actual grasp success rates. We additionally investigate the effect of adding the missing two elements – object dynamics and pose uncertainty – one by one or at the same time to the grasping simulation, and see which method is most effective in estimating the actual grasp success or failure, which is the core of this study. Eight combinations of the simulation elements have been tested for this purpose as listed below[7]:

- **Sa:** Static object + Measure A *(Existing method)*
  In this method the fingers close kinematically until they touch the object or reach joint limits and the object is assumed to be static during the grasping. Then, the magnitude of the maximum wrench that the grasp can resist is measured by calculating the minimum distance to the boundary of the convex hull of the contact wrenches (Ferrari and Canny [11]).
- **SaU:** Static object + Measure A + Uncertainty
  We added the object pose uncertainty to the existing method by running multiple simulations for each grasp and averaging the quality scores as described in Section III-C.
- **Da:** Dynamic object + Measure A
  We run the physically-based grasping simulation described in Section III-A, and evaluate the grasp quality with the existing force-closure based measure.
- **DaU:** Dynamic object + Measure A + Uncertainty
  The uncertainty in the object pose is added to the previous method by running multiple simulations for each grasp.
- **Db:** Dynamic object + Measure B
  After running the physically-based grasping simulation, we evaluate the quality of a grasp using the 3-step scoring system based on the number of contact links at the final grasp configuration.
- **DbU:** Dynamic object + Measure B + Uncertainty
  The pose uncertainty is added to the previous method by running the simulation multiple times for each grasp.
- **Dc:** Dynamic object + Measure C
  After running the physically-based grasping simulation, we calculate the score values using (2) from the translational and rotational deviations at the final grasp configuration. We set the deviation limit L to 5 cm and 30 deg for position and orientation respectively.
- **DcU:** Dynamic object + Measure C + Uncertainty
  Again, the uncertainty in the object pose is added to the previous method by running multiple simlations for each grasp.

For each object, we computed the quality scores of the same 10 grasps, which we had tested in the experiments, using the simulation methods listed above. Then, we evaluated the ability of the methods to predict the actual grasp success or failure by calculating the correlation coefficient between the simulated scores and the experimental data (visual inspection), which is summarized in Figure 6.

In the top row of Figure 6 we compare the grasp quality scores from the existing method (Sa) with the actual grasp success rates obtained from the experiment. The numbers in gray denote the correlation coefficients between the two score sets. As seen from the graphs, the two data sets for each object have no consistency and, as a result, the correlation is very low. This implies that the existing method does not predict the actual grasp success or failure well. In the paragraphs below we explore the ways of possible improvements listed above by adding missing factors, such as pose uncertainty and dynamics, to the simulation or applying different quality measures, and investigating their effects on the performance.

[6]We use Pearson's correlation defined as $\rho = \frac{\mathrm{cov}(X,Y)}{\sigma_X \sigma_Y}$ where cov means covariance and σ_X and σ_Y are the standard deviations of the two data sets X and Y.

[7]The label characters 'S', 'D', 'U', 'a', 'b' and 'c' represent 'static', 'dynamic', 'uncertainty', 'Measure A', 'Measure B', and 'Measure C' respectively. See Section III-B for the description of the measures.

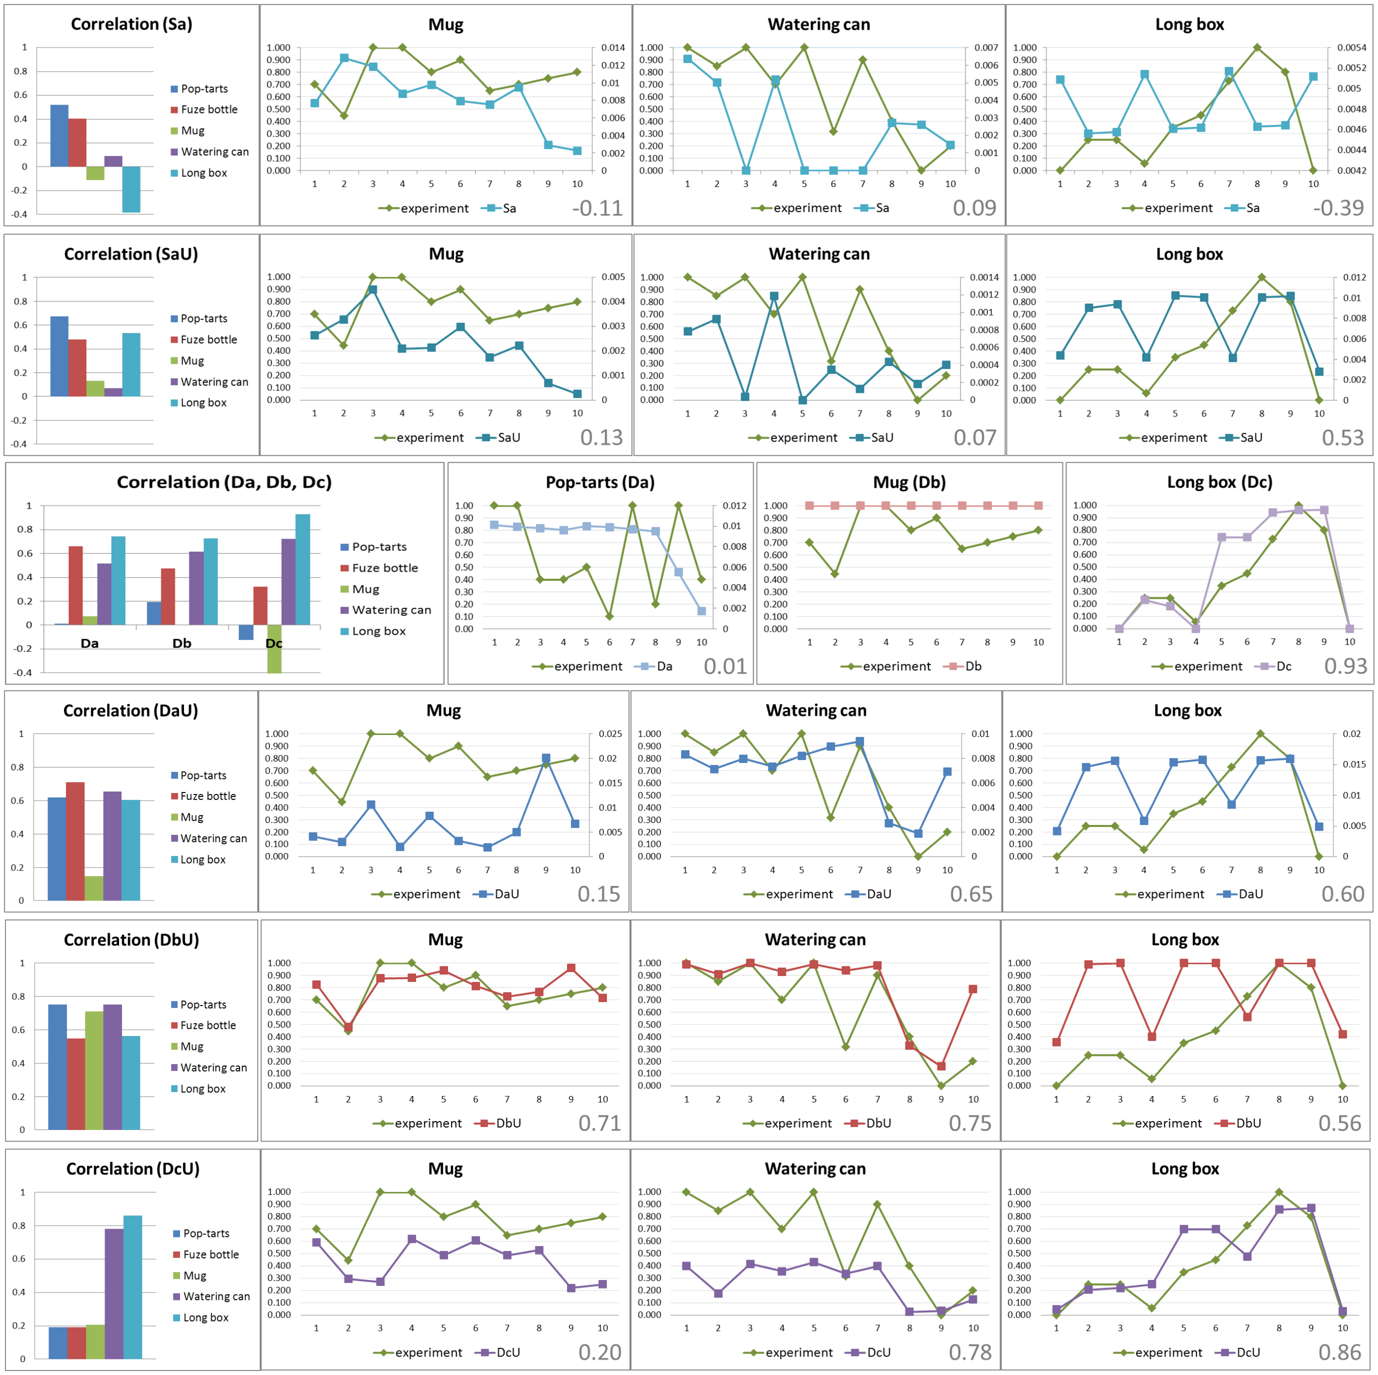

Fig. 6. The performance of the simulation-based methods in predicting the actual grasp success rates obtained from the experiments. **(Left)** The correlation coefficients of the tested methods. **(Right)** Comparison of the quality scores from the methods and the experiments. The number in gray at the lower right corner is the corresponding correlation coefficient. (Correlation is not defined for Mug(Db) because standard deviation of the Db scores is zero.)

We first tried adding pose uncertainty to the existing method by repeating the kinematic grasping where the initial object pose relative to the hand is set slightly differently by sampling the pose error from a normal distribution model. In our test, however, this method (SaU) did not make meaningful improvement in estimating the actual grasp success rates as shown in the second row of Figure 6. For some objects, the correlation coefficients have increased compared to the existing method, but overall, the correlations are not high enough to be effectively used for estimating the likely success or failure of actual robot grasping.

The third row of Figure 6 shows the results of considering object dynamics in the grasping simulation. The three measures described in Section III-B were used to evaluate the grasp quality after grasping and their corresponding results were marked as Da, Db, and Dc respectively. The pose uncertainty was not considered here. Overall, the methods did not show consistency across the objects – for example, the methods Db and Dc showed relatively good correlation for the two big objects (Watering can and Long box), but not for the other

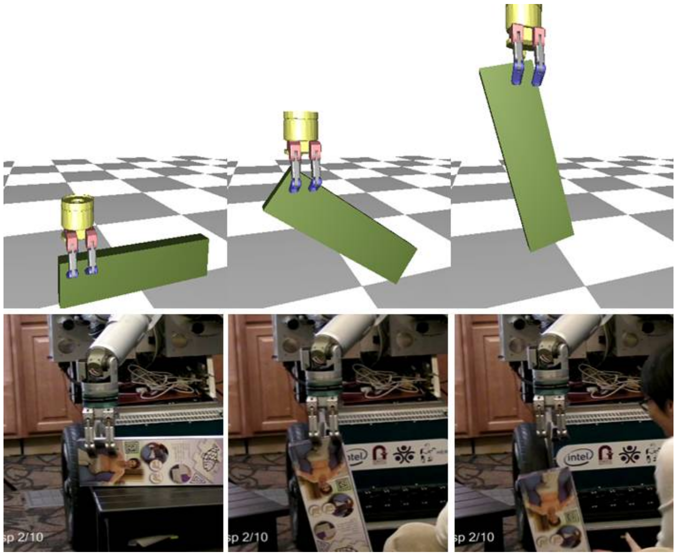

Fig. 7. Lifting up Long box in simulation (upper) and experiment (lower). In the simulation the hand was able to hold the grip of the object even after the rotation by the gravity, but in the experiment, the robot hand dropped the object in such a situation. In our experiment with Long box, instead of having the robot lift it up automatically, we had to manually remove the base support carefully due to the limitation of the robot arm workspace.

objects. Especially, the quality score from the method Db tends to be all-or-nothing and this makes it difficult to discriminate better grasps effectively. The third graph (Mug) in the third row of the figure is a good example of this – all of the 10 grasps got the same score of 1 because they ended up with successful grasps with three or more contact points, which could happen quite often in the simulation even though it may not be so plausible in the real world. Thus, in such situations, we anticipate that adding pose uncertainty to the simulation would be effective in discriminating better grasps from worse ones.

In order to investigate the effect of considering object dynamics and pose uncertainty at the same time, we have tested three methods (DaU, DbU, and DcU) where each of them uses different measure to evaluate the quality of a grasp while using the same physically-based grasping simulation and pose uncertainty model. In our experiment the method DbU, using the measure based on the number of contact links (Measure B), showed the best consistency with good performance in estimating the actual grasp success rates of the tested objects (Figure 6, fifth row). This is perhaps because the method uses the same 3-step scoring system based on the number of contact links which has also been used in the experiments.

The method DbU, however, showed some limitation in distinguishing better performing grasps from others for some objects. For example, in our test, the method showed a poor performance compared to our initial expectation for the long box. We had anticipated that the method would work very well for the object because the movement of the long object during grasping is largely affected by the gravity and our physically based simulation can effectively capture this. It turned out that the relatively poor performance was because the slip movements after the large rotation occurring in the lift-up stage were not correctly captured by our simple frictional contact mechanism based on a penalty-based method and the Coulomb friction model. More specifically, as shown in Figure 7, our grasping simulation tends to hold the long box even after the large rotation by the gravity while the real robot hand dropped the object quite often in such a situation, and this is why counting the number of contacts did not work well for the object.

Using the force-closure based measure (Measure A) in the method DaU also resulted in a relatively good consistency across the objects except for the small and light plastic mug (Figure 6, fourth row). It is very difficult to reproduce the complicated actual movements and the final grasp configuration through a physically-based simulation for such a small and light rigid object due to the limitation in model correctness that can be achieved in practical time and effort. Thus, the simulated contact points at the final grasp configuration can be quite different from the actual contact locations seen in the experiment. For this reason, we speculate that the force-closure based measure can be more sensitive to the simulation error than the simple measuring mechanism counting the number of contacts after grasping (Measure B), especially when both were applied to predicting the likely success or failure of a grasp.

We did not consider the pose deviation in measuring the grasp quality in our experiments, so it would not be so surprising to see the method DcU, measuring the pose deviation at the final grasp configuration, show very low correlations to the experimental data (Figure 6, bottom row). Interestingly, however, the method worked for the large objects (Watering can and Long box). We speculate this is because a large pose deviation of the objects during the grasping and the lift-up stages usually resulted in the loss of grip in our experiments, and such a situation is handled by penalizing the pose deviation in the method DcU.

In summary, the existing method (Sa) showed poor performance in estimating the likely success or failure of the actual robot grasping. Adding pose uncertainty to the existing method (SaU) increased the correlation to the experimental data slightly, but still the performance is not enough for using the method in predicting the actual grasp success rates. Considering object dynamics (Da, Db, and Dc) can make some improvement in simulating the actual robot grasping, but their performance was not consistent across the objects. Finally, we observed that considering both object dynamics and pose uncertainty in the simulation can bring a significant improvement in performance with consistency as shown in the method DbU.

B. Sensitivity Testing

We have used the same setting for pose uncertainty in every simulation result shown in Section V-A, but the actual uncertainty may vary depending on many factors such as the object type and even the location of the object as mentioned in Section III-C. In order to investigate the effect of parameter change on the simulation results, we repeated the

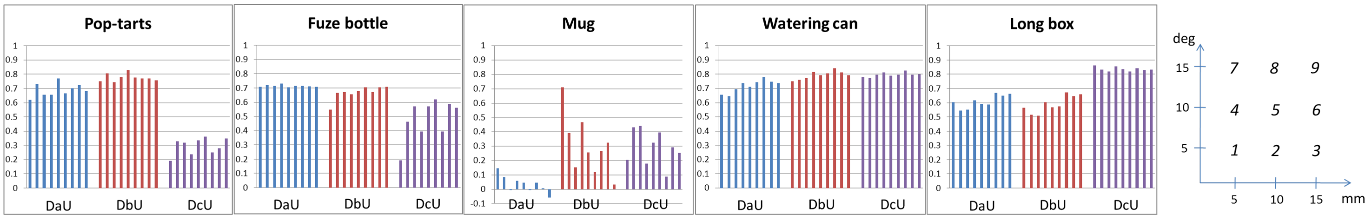

Fig. 8. Correlation for various parameters for pose uncertainty. The blue, red, and magenta bar clusters represent the method DaU, DbU, and DcU respectively. For each method, we tested nine sets of parameters for pose uncertainty. The rightmost diagram shows the parameter values we used in the test and their numbering – for example, in order to obtain the first bar (from left) of each method, we set $(e_p, e_R) = (5\text{mm}, 5\text{deg})$ as the expectations of the half-normal distributions for the positional and rotational errors, which is the original setting of the simulations shown in Section V-A.

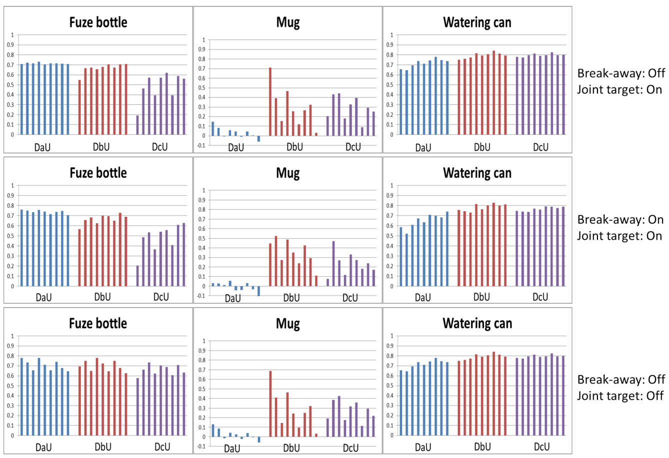

Fig. 9. Correlation for various settings for finger closing. Each row represents the result of a different finger closing setting where the top row is the result of the original setting used in Section V-A. The bar clusters in blue, red and magenta represent the methods DaU, DbU, and DcU, which are considering object dynamics and pose uncertainty, and each bar in the clusters shows the correlation value of a particular uncertainty setting described in Figure 8.

same analysis multiple times with different sets of uncertainty model parameters and showed how the performance of the methods changes depending on the parameter setting in Figure 8. Though the individual score values are somewhat different depending on the parameters, overall, the performance (or correlation) of each method was well preserved in most of the objects except for the plastic mug. This suggests that, as long as the object is not too small and light, considering pose uncertainty with a reasonable parameter setting would be still beneficial to predicting the likely success or failure of the actual grasping more correctly, even though we do not have good information on the actual uncertainty. However, in case of the small and light object (Mug), the simulation result was very sensitive to the change of parameters. We speculate that this is because the small size and light weight increases uncertainty in the dynamic response of the object to finger closing, and as a result, a less accurate uncertainty model can degrade the performance of the methods more significantly.

We have also investigated the effect of change in our simulation setting for finger closing such as break-away and target joint angles. In the simulations in Section V-A we disabled the break-away function and set the target joint angles same as the ones used in the experiments. Two additional settings have been tested – first, the break-away function was enabled, and second, the target joint angles were disabled so that the fingers can close further up to their joint limits – and the results are shown in Figure 9. Changing the break-away condition affects interaction between the fingers and the object during grasping directly, so this caused some change in the simulation results.[8] As expected, the light and small object (Mug) was more affected by this than the other objects. Closing the fingers further, however, did not make meaningful change for most objects. Only Fuze bottle showed some change in the correlation coefficients because the power grasping around the narrow neck of the bottle object can be directly affected by the change in the target joint angles. Note that, however, the overall trend in the correlation coefficients across the methods and the uncertainty settings was well preserved in both cases, and this also supports our finding that considering both object dynamics and pose uncertainty improves the performance in predicting the actual grasp success rates consistently even with the presence of a small range of errors in modeling.

TABLE II
FRICTION COEFFICIENTS USED IN OUR SIMULATION

Objects	Friction coefficients (μ_s, μ_d) (finger tip)	(table)
Pop-tarts	0.50, 0.30	0.30, 0.20
Fuze bottle	0.50, 0.30	0.30, 0.20
Mug	0.47, 0.39	0.19, 0.14
Watering can	0.40, 0.30	0.23, 0.14
Long box	1.53, 0.45	0.36, 0.19

μ_s : static coefficients, μ_d : dynamic coefficients

C. Simulation Setting

Considering object dynamics in the grasping simulation is very computationally expensive compared to the kinematic simulation used in the existing method because the step size for integration must become much smaller in order to capture the fast and complicated response of the object by the collision during grasping and suppress numerical instability. In our dynamic simulation we set the step size as 0.5 msec for the simple objects such as Pop-tarts, and 0.2 msec for more complex objects such as Mug, and this requires up to 20,000 iterations of collision checking and solving the equations of motion of the rigid object for a single grasping simulation.

[8]This is just for investigating the effect of closing mechanism change on the simulation result. The breakaway cannot be disabled in the actual Barrett hand used in our experiment.

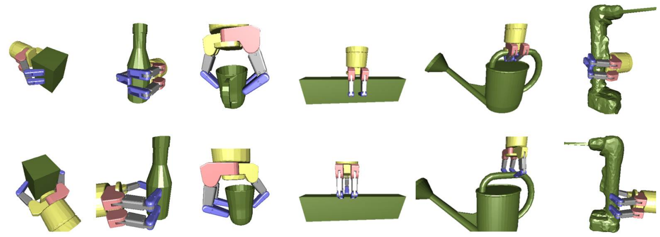

Fig. 10. Representative grasps chosen by the method considering dynamics and uncertainty (DcU+DbU, upper) and the existing method (Sa, lower). See Figure 11 for the top 10 grasps of each object using each of the two metrics.

In our implementation, it took about 3.6 sec in average to physically simulate Mug grasping on a desktop PC with an Intel Core 2 CPU running at 2.83 GHz where more than 90% of the computation time was spent for collision checking. On the other hand, in the kinematic grasping for the existing method, at most about 200 iterations of coarse and fine stepping suffice to obtain the contacts between the object and the robot hand. It took about 0.1 sec in average to run the kinematic grasping with OpenRAVE on the same machine.

We used a penalty-based contact model by Yamane and Nakamura [34] and Coulomb friction model. The friction coefficients were obtained by measuring the forces that we applied to push the object while having the robot hand grab it with a pinch grasp, and the grasping force of the hand.[9] Table II shows the friction coefficients we used in our simulation.

VI. Application: Grasp Set Generation

We have applied the simulation-based evaluation methods to automatically generating grasp sets. For each object, 1000 grasps were randomly chosen from a myriad number of grasp candidates which were generated using a simple geometry-based sampling technique [3] implemented in OpenRAVE. We set positive values for the distance between the palm and the object in sampling to avoid unwanted collision at the pre-grasp stage. Then, the quality scores of each grasp candidate were evaluated using the methods. Note that the grasp evaluation methods are independent of the sampling method, so any sampling technique for generating grasp candidates can replace the current sampler.

In Figure 10 we compared the typical grasps chosen by the method considering dynamics and uncertainty (DcU+DbU), and by the existing method (Sa). Also, see Figure 11 for the top 10 grasps of each object using each of these two metrics. For DcU+DbU, we chose the grasps with minimal pose deviation (Measure C, DcU) among the grasps having more than 90% of simulated grasping success rates (Measure B, DbU).[10]Note that any combination of the measures can be used here. For example, one can add DaU to this in order to ensure the final grasp configuration forms a force-closure which is not a feature of DbU and DcU. As seen in the figures, the method based on dynamic simulation tends to choose power grasps using both palm and fingers for better success rates, while the existing method sticks to the pinch grasps as mentioned in Section I. We will discuss this difference further later in this section.

In Figure 12 we showed the top 30 grasps chosen by Sa, SaU, and DcU+DbU respectively. Obviously, the existing method (Sa) does not have an ability to recognize unrealistic grasps such as edge grasp which is likely to cause a failure in the real environment. Adding uncertainty to the existing method (SaU) can improve this to some extent, but still there remain many unrealistic grasps in the chosen grasp set, while most of such grasps were excluded by considering dynamics and uncertainty (DcU+DbU). In Figure 13 the chosen grasps are marked as blue circles (Sa), green squares (SaU), and red dots (DcU+DbU) in the scatter plot where the dots represents the simulated quality scores (DbU and DcU) of the 1000 grasp candidates. In the simulation test, 10 out of the top 30 grasps had a simulated grasp success rate less than 40% for both Sa and SaU and the average rate of the 30 grasps was about 60% for both, which indicates that the grasp sets generated with the existing method and its variant considering uncertainty are likely to cause frequent grasp failure in the real robot grasping.

Interestingly, the method considering both dynamics and uncertainty (DcU+DbU) chose quite realistic grasps in type and location that are very similar to the actual human grasps on the handle (Figure 12, right). We could also observe a similar result for the drill object, and showed the accumulated finger prints of the grasps in Figure 14. Apparently, if available, the method based on dynamic simulation tends to choose a grasp utilizing the design feature for grasping such as handles to increase the grasp success rates, which results in more robust and natural grasps.

In order to consider the pose uncertainty in the grasp quality evaluation, we ran the grasping simulation 100 times for each grasp. This means we had to run 100,000 iterations of the grasping simulation to choose good grasps from the 1000 grasp candidates. This took about 4 days for the dynamic simulations and about 3 hours for the kinematic simulations in case of the plastic mug. The long computation time is not so critical in offline running for known objects, but prevents the method from being used for unknown objects online. Note that, however, the grasping simulations are independent of each other, and thus parallelizable.

The computation time can be significantly reduced by using heuristic grasp filtering. In our experience, if the robot hand has failed to grab an object located at its nominal position (without uncertainty) with a grasp, the grasp is not likely to become a good grasp and we do not need to test the grasp any more. Also, if the averaged grasp success rate from the first 10 simulations (with uncertainty) is lower than a reference value (90% in our test), we can discard the grasp instead of proceeding the remaining 90 simulations. Figure 15 shows an example of heuristic grasp filtering for the watering can where the green dots and the pink dots represent the grasps that can be discarded after the first and the tenth simulation

[9]For some objects, we failed in measuring the friction coefficients and used moderately reasonable values instead. For example, in the case of Watering can, the trunk was too big and the handle was too narrow to be grabbed by the finger tips for measuring friction. For Fuze bottle, we failed to get reasonable force data during the simple friction test possibly because of the local deformation on the plastic bottle surface by the pinch grasp.

[10]See, for example, the red dots in Figure 13.

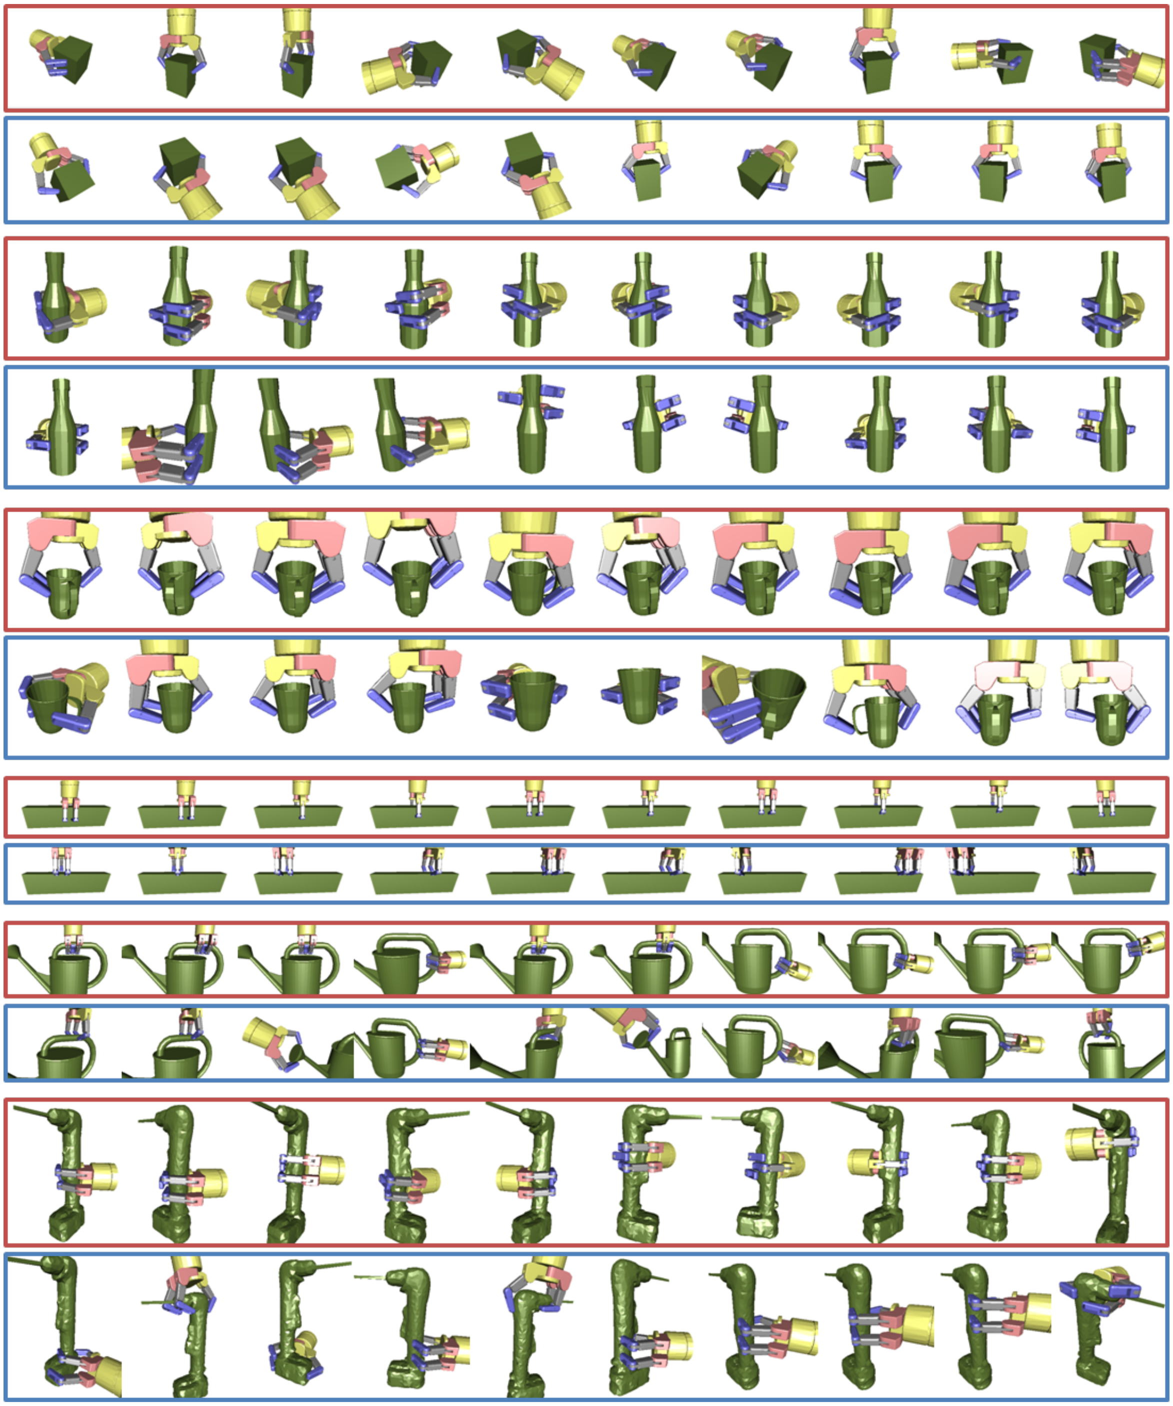

Fig. 11. The top 10 grasps chosen by DcU+DbU (upper, red boxes) and Sa (lower, blue boxes). Simulated final configurations of the grasps are shown here for convenience in visual comparison.

respectively. Less than 200 (shown as the black dots) out of the 1000 grasp candidates need the full 100 simulations to consider the pose uncertainty, which can save up to 80% of the computation time.

As mentioned above, any sampling technique for generating grasp candidates can be incorporated with the simulation-based grasp evaluation methods. In particular, running dynamic simulations multiple times for every grasp candidate is a very expensive process, so reducing the sampling search space by choosing an appropriate method can greatly improve overall processing speed. For example, one can use other sampling techniques such as eigengrasps (a low-dimensional basis for grasp configuration) [8], a superquadratic decomposition tree [14], and a grid of medial spheres [28] to reduce the search space significantly without sacrificing potentially high quality grasp candidates too much. We speculate that the existing force-closure based method can also benefit from such a sampling technique by taking more meaningful and natural grasp candidates into account. Note that, however, even in such a case, our evaluation method based on dynamic simulation

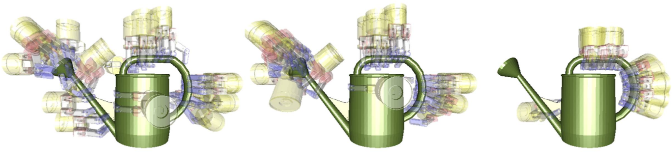

Fig. 12. The top 30 grasps chosen by the existing method (Sa, left), the existing method with uncertainty (SaU, middle), and the method considering both dynamics and uncertainty (DcU+DbU, right).

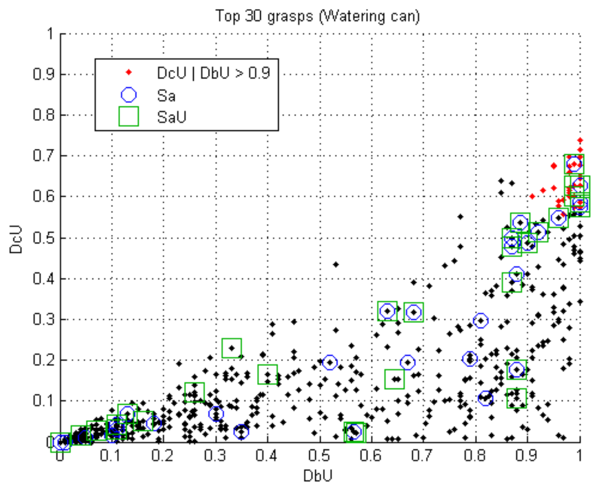

Fig. 13. The simulated quality scores (DbU and DcU) of the 1000 grasp candidates for the watering can. The blue circles, green squares, and red dots represent the top 30 grasps chosen by Sa, SaU, and DcU+DbU respectively.

would still work better than the existing force-closure based method in predicting the likely success or failure of a grasp candidate (Section V) and thus build a better grasp set by choosing more robust grasps among the candidates.

Our dynamic simulation based method tends to favor power grasps while the existing method (Sa) seems to prefer pinch grasps (Figure 10 and 11). We should first note that the reason why the existing method prefers pinch grasps is not because the force-closure based quality measure (Measure A) favors pinch grasps, but because, as mentioned above, we set positive values for the distance between the palm and the object in grasp sampling to avoid unwanted collision at the pre-grasp

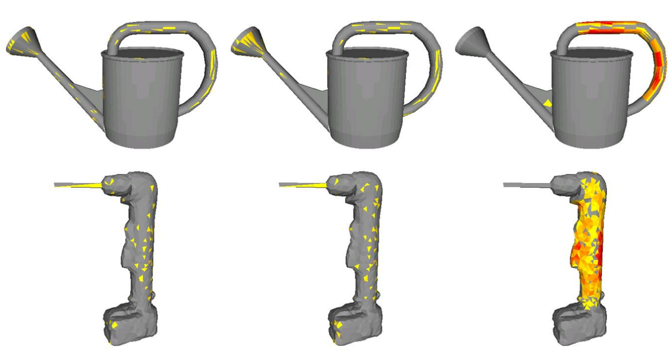

Fig. 14. Accumulated finger/hand prints of the top 30 grasps chosen by Sa (left), SaU (middle), and DcU+DbU (right). The color varying from yellow to red represents the density of the prints.

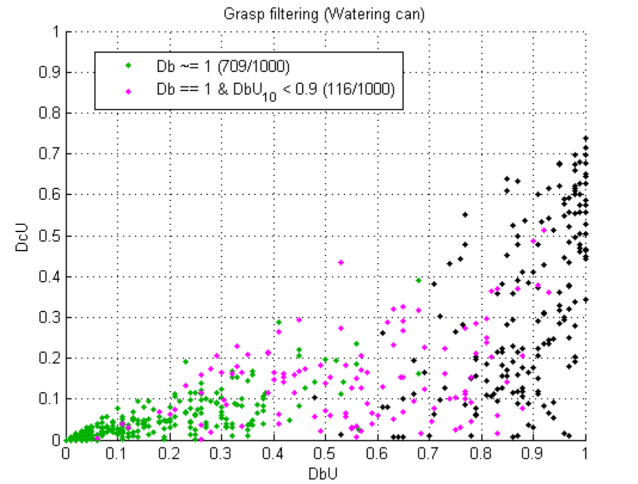

Fig. 15. An example of heuristic grasp filtering for saving the massive simulation time in considering uncertainty. The green dots and the pink dots represent the grasps that can be discarded after the first and the tenth simulation respectively without proceeding the remaining 99 and 90 simulations.

stage and the employed kinematic finger closing with the static object assumption cannot create contacts on the palm unless the palm was already in contact with the object before closing the fingers. Thus, most of the power grasp candidates were discarded because they could not form a force-closure grasp due to lack of palm contacts.

One simple way to include power grasps would be to allow initial contact at the palm by setting the distance from the palm to the object to zero in sampling. This sampling strategy has also been used in a power grasp planner [29]. However, the grasps are likely to cause unwanted collision at the pre-grasp stage, and thus a more sophisticated planning and control mechanism would be needed to reach the generated grasps carefully without pushing away the object. One alternative to this is to keep setting positive palm distance in sampling, as we did in this study, and use a grasp quality metric that can handle pre-grasp configurations without contacts. The quality metric used by Ciocarlie and Allen [8] is a good example of this. They assess the quality of a grasp, where the fingers are not in contact yet, by assuming the hand can apply potential contact wrenches at pre-determined desired contact locations on fingers and palm. The potential contact wrenches are scaled depending on the distances between the desired contact locations and the object, and thus the contact locations closer to the object make larger contribution to the grasp wrench space.

Dynamic simulation, on the contrary, can create contacts on the palm naturally by pulling the object inward toward the palm with the closing fingers. The object wrapping with whole hand can work quite robustly even under the pose uncertainty, and this gives power grasps higher scores or higher grasp success rates. Thus, the grasp set generator, which picks grasps with high scores, is more likely to choose power grasps than pinch grasps.

Note that, however, the grasp quality measures (Measure A, B and C) can be applied to any set of grasp candidates. This means that, if a particular type of grasp set is needed, one can first pick grasps satisfying the requirement using an appropriate quality measure, and then examine the grasps further with our dynamic simulation based measure to pick more robust ones among them. For example, if a pinch grasp set is needed to manipulate an object, the task oriented quality measure by Li and Sastry [22] in conjunction with the kinematic grasping simulation can be used to filter out unwanted grasp candidates, and then choose grasps with better grasp success rates among them using our dynamic simulation based method. The same strategy can also be applied to refining an existing grasp set database such as [15].

VII. Conclusion

Evaluating the quality of a grasp correctly is necessary to improve the robustness of robot grasping. We investigated the performance of several simulation-based algorithms in predicting the likely success or failure of a grasp in the real environment through experiments. We observed that the force-closure based existing method has poor performance and this can be significantly improved by considering object dynamics and pose uncertainty at the same time in the evaluation. The new method considering both dynamics and uncertainty was able to estimate the actual grasp success rates more correctly and with more consistency than the existing method and its possible variant with added uncertainty.

We have also applied the new algorithm to generating grasp sets automatically and were able to obtain grasp sets with better quality than the existing method. The new method tends to choose, if available, stable and realistic grasps such as the power grasps on the handle. We anticipate that, when applied to a fully automated robot grasping planner, such a grasp set would improve not only the robustness but also the naturalness of robot grasping.

Most importantly, we showed that dynamic simulation, which is often thought to be difficult to apply to an evaluation problem involving rapidly changing contacts, can be effectively used to improve the performance of estimating the actual grasp success rate of a grasp in the real environment. The straightforward approximation we used in our physically based grasping simulation (e.g., the kinematics-based hand closing mechanism) was quite effective in capturing the effect of object dynamics in open-loop grasping with minimal effort and time for modeling and implementation. Running the low cost simulation multiple times to consider uncertainty significantly improved the performance of the simulation based algorithm with consistency even without having precise models for, e.g., the robot hand and pose uncertainty.

Although promising, the presented method has limitations that require further improvement in the future. We have shown object dynamics and pose uncertainty are key important factors, but there are many other factors that can significantly influence grasping in the real environment. For example, we expect capturing frictional contact more correctly with a sophisticated model would improve the accuracy of simulation of the subtle and complex phenomena in slip contact and lead to a better performance in predicting the grasp success rates. Fully considering hand dynamics and control loop would also improve the performance, though building a precise model for a robotic hand system requires a great effort and time in general. It would be interesting to extend the grasp quality evaluation method, currently applied to open-loop grasping, to more sophisticated grasping control mechanisms such as force compliant grasping [17]. Data portability is also an important issue in practice. For example, it would become an interesting topic for future research to investigate how well a grasp set generated assuming one grasping mechanism would work with another grasp controller.

Acknowledgment

The authors would like to thank Professor Siddhartha Srinivasa and his lab members at Carnegie Mellon University for allowing and helping them to use HERB in their experiments. This work has been supported by Toyota Motor Corporation (award #409985) and NSF awards IIS-1145640 and IIS-1218182.

References

[1] Bullet: Game physics simulation.

[2] R. Balasubramanian, L. Xu, P.D. Brook, J.R. Smith, and Y. Matsuoka. Human-guided grasp measures improve grasp robustness on physical robot. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 2294–2301. IEEE, 2010.

[3] Dmitry Berenson, Rosen Diankov, Koichi Nishiwaki, Satoshi Kagami, and James Kuffner. Grasp planning in complex scenes. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids07)*, December 2007.

[4] R.C. Brost. *Planning robot grasping motions in the presence of uncertainty*. Carnegie Mellon University, The Robotics Institute, 1985.

[5] E. Chinellato, A. Morales, R.B. Fisher, and A.P. del Pobil. Visual quality measures for characterizing planar robot grasps. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 35(1):30–41, 2005.

[6] V.N. Christopoulos and P. Schrater. Handling shape and contact location uncertainty in grasping two-dimensional planar objects. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, pages 1557–1563. IEEE, 2007.

[7] M. Ciocarlie, A. Miller, and P. Allen. Grasp analysis using deformable fingers. In *Intelligent Robots and Systems, 2005. (IROS 2005). 2005 IEEE/RSJ International Conference on*, pages 4122 – 4128, aug. 2005.

[8] M.T. Ciocarlie and P.K. Allen. Hand posture subspaces for dexterous robotic grasping. *The International Journal of Robotics Research*, 28(7):851–867, 2009.

[9] Alvaro Collet, Manuel Martinez, and Siddhartha S. Srinivasa. The MOPED framework: Object Recognition and Pose Estimation for Manipulation. 2011.

[10] Rosen Diankov. *Automated Construction of Robotic Manipulation Programs*. PhD thesis, Carnegie Mellon University, Robotics Institute, August 2010.

[11] C. Ferrari and J. Canny. Planning optimal grasps. In *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, pages 2290–2295. IEEE, 1992.

[12] K.Y. Goldberg and M.T. Mason. Bayesian grasping. In *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*, pages 1264–1269. IEEE, 1990.

[13] C. Goldfeder and P.K. Allen. Data-driven grasping. *Autonomous Robots*, pages 1–20, 2011.

[14] C. Goldfeder, P.K. Allen, C. Lackner, and R. Pelossof. Grasp planning via decomposition trees. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 4679 –4684, april 2007.

[15] C. Goldfeder, M. Ciocarlie, H. Dang, and P.K. Allen. The columbia grasp database. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, pages 1710–1716. IEEE, 2009.

[16] Kaijen Hsiao, Tomas Lozano-Perez, and Leslie Pack Kaelbling. Relatively robust grasping. In *ICAPS 2009 Workshop on Bridging the Gap Between Task and Motion Planning*, 2009.

[17] Moslem Kazemi, Jean-Sebastien Valois, J. Andrew (Drew) Bagnell, and Nancy Pollard. Robust object grasping using force compliant motion primitives. In *Robotics: Science and Systems*, July 2012.

[18] B.H. Kim, S.R. Oh, B.J. Yi, and I.H. Suh. Optimal grasping based on non-dimensionalized performance indices. In *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, volume 2, pages 949–956. IEEE, 2001.

[19] Junggon Kim, Kunihiro Iwamoto, James J. Kuffner, Yasuhiro Ota, and Nancy S. Pollard. Physically-based grasp quality evaluation under uncertainty. In *IEEE International Conference on Robotics and Automation*, 2012.

[20] J.J. Kuffner Jr and S.M. LaValle. Rrt-connect: An efficient approach to single-query path planning. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, volume 2, pages 995–1001. IEEE, 2000.

[21] Beatriz Len, Stefan Ulbrich, Rosen Diankov, Gustavo Puche, Markus Przybylski, Antonio Morales, Tamim Asfour, Sami Moisio, Jeannette Bohg, James Kuffner, and Rdiger Dillmann. Opengrasp: A toolkit for robot grasping simulation. In *Simulation, Modeling, and Programming for Autonomous Robots*, volume 6472 of *Lecture Notes in Computer Science*, pages 109–120. Springer Berlin / Heidelberg, 2010.

[22] Z. Li and S.S. Sastry. Task-oriented optimal grasping by multifingered robot hands. *Robotics and Automation, IEEE Journal of*, 4(1):32–44, 1988.

[23] A. Miller, P. Allen, V. Santos, and F. Valero-Cuevas. From robotic hands to human hands: a visualization and simulation engine for grasping research. *Industrial Robot: An International Journal*, 32(1):55–63, 2005.

[24] A.T. Miller and P.K. Allen. Graspit! a versatile simulator for robotic grasping. *Robotics & Automation Magazine, IEEE*, 11(4):110–122, 2004.

[25] B. Mirtich and J. Canny. Easily computable optimum grasps in 2-d and 3-d. In *Robotics and Automation, 1994. Proceedings., 1994 IEEE International Conference on*, pages 739–747. IEEE, 1994.

[26] A. Morales, E. Chinellato, A.H. Fagg, and A.P. del Pobil. Experimental prediction of the performance of grasp tasks from visual features. In *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 4, pages 3423 – 3428 vol.3, oct. 2003.

[27] Y.C. Park and G.P. Starr. Grasp synthesis of polygonal objects using a three-fingered robot hand. *The International journal of robotics research*, 11(3):163–184, 1992.

[28] Markus Przybylski, Tamim Asfour, and Rudiger Dillmann. Planning grasps for robotic hands using a novel object representation based on the medial axis transform. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 1781 –1788, sept. 2011.

[29] M.A. Roa, M.J. Argus, D. Leidner, C. Borst, and G. Hirzinger. Power grasp planning for anthropomorphic robot hands. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 563–569, 2012.

[30] Russell Smith. Ode: Open dynamics engine.

[31] S.S. Srinivasa, D. Ferguson, C.J. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M.V. Weghe. Herb: a home exploring robotic butler. *Autonomous Robots*, 28(1):5–20, 2010.

[32] R. Suárez Feijóo, J. Cornellà Medrano, M. Roa Garzón, et al. Grasp quality measures. Technical Report IOC-DT-P-2006-10, Universitat Politecnica de Catalunya, 2006.

[33] J. Weisz and P.K. Allen. Pose error robust grasping from contact wrench space metrics. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 557–562. IEEE, 2012.

[34] K. Yamane and Y. Nakamura. Stable penalty-based model of frictional contacts. In *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 1904–1909. IEEE, 2006.

[35] L.E. Zhang, J. Betz, and J.C. Trinkle. Comparison of simulated and experimental grasping actions in the plane. In *First Joint International Conference on Multibody System Dynamics*, May 2010.

[36] Y. Zheng and W.H. Qian. Coping with the grasping uncertainties in force-closure analysis. *The International Journal of Robotics Research*, 24(4):311–327, 2005.