Physically-based Grasp Quality Evaluation under Uncertainty*

Junggon Kim[1], Kunihiro Iwamoto[2], James J. Kuffner[3], Yasuhiro Ota[2], and Nancy S. Pollard[1]

***Abstract*— In this paper new grasp quality measures considering both object dynamics and pose uncertainty are proposed. Dynamics of the object is incorporated into our grasping simulation to capture the change of its pose and contact points during grasping. Pose uncertainty is considered by running multiple simulations starting from slightly different initial poses sampled from a probability distribution model. A simple robotic grasping strategy is simulated and the quality score of the resulting grasp is evaluated from the simulation result. The effectiveness of the new quality measures on predicting the actual grasp success rate is shown through a real robot experiment.**

I. INTRODUCTION

Grasping objects with mechanical hands stably and reliably remains a challenging task. In this paper we estimate the quality of robotic grasping under uncertainty using a simulation based algorithm, and compare the results with experiments. Our grasp quality measure uses a physically based simulation to consider the complicated interaction between the robotic hand and an object that occurs when grasping. The uncertainty in the simulation input data such as the object position obtained from sensors is considered in the algorithm by sampling from a probability distribution model and running multiple simulations for each grasp.

We focus on estimating the success rate of an open-loop grasp defined as a relative pose of the hand to the object and finger configuration prior to grasping. Better estimation of likely success or failure of such an open-loop grasp is useful in itself for systems that rely on a relatively simple finger closing mechanism for grasping. It can also be used in conjunction with more sophisticated feedback driven grasping algorithms to remove poorly performing grasps from consideration.

Many planning algorithms for robot manipulation use a set of precomputed grasps in order to pick a reachable grasp and plan a trajectory to it within a reasonable time. A large grasp set is usually used to increase the possibility of finding a feasible grasp solution which is reachable from the current robot configuration. Thus it is necessary to collect a large selection of 'good' grasps because the success or failure of the manipulation task is directly affected by the quality of the grasp set used in the planning.

*This work has been supported by Toyota Motor Corporation, Award #409985.

[1]Junggon Kim and Nancy Pollard are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. {junggon,nsp}@cs.cmu.edu

[2]Kunihiro Iwamoto and Yasuhiro Ota are with the Partner Robot and Advanced Engineering Group, Toyota Motor Engineering and Manufacturing North America, CA 94043, USA. {kunihiro.iwamoto,yasuhiro.ota}@tema.toyota.com

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

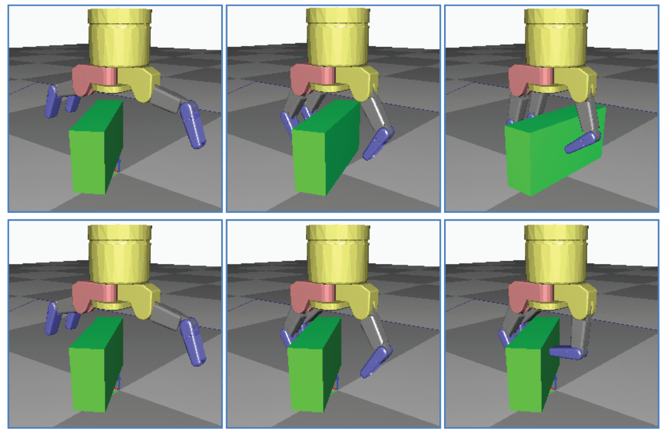

Fig. 1. Dynamic (upper) and static (lower) objects. The dynamic object can translate and rotate according to the finger motion, which results in more plausible contact points than the static object.

Checking force-closure from contacts is a simple but useful way to check if a grasp is good or stable enough to resist external disturbance forces, and this approach has been used in many automatic grasp generation tools such as GraspIt! to choose good grasps among myriad numbers of possible grasps. Various methods have been suggested to qualitatively measure how good a grasp is, and those quality measures are often used in sorting the grasps so that better grasps can be more likely to be used in applications such as manipulation planning.

To compute the quality score of a grasp with existing methods, we need to know about the contacts between the hand and the object. A kinematic simulation which closes the fingers until touching the object is often used to obtain the contact points under the assumption that the object remains at the same place. However, such a static object assumption does not hold in many grasping situations as shown in Fig. 1. When the resulting grasp is far from that obtained using the static object assumption, contact information will be incorrect, as will the estimate of grasp quality.

Uncertainty in the data obtained from sensors can also affect the success rate of a proposed grasp. For example, a slight change in object position can change a success into a failure, or vice versa, as shown in Fig. 2. Such an abrupt change in grasping result can only be effectively captured with a dynamic grasping simulation as presented here.

In this paper, we incorporate the dynamics of the object into our grasping simulation to capture the change of its

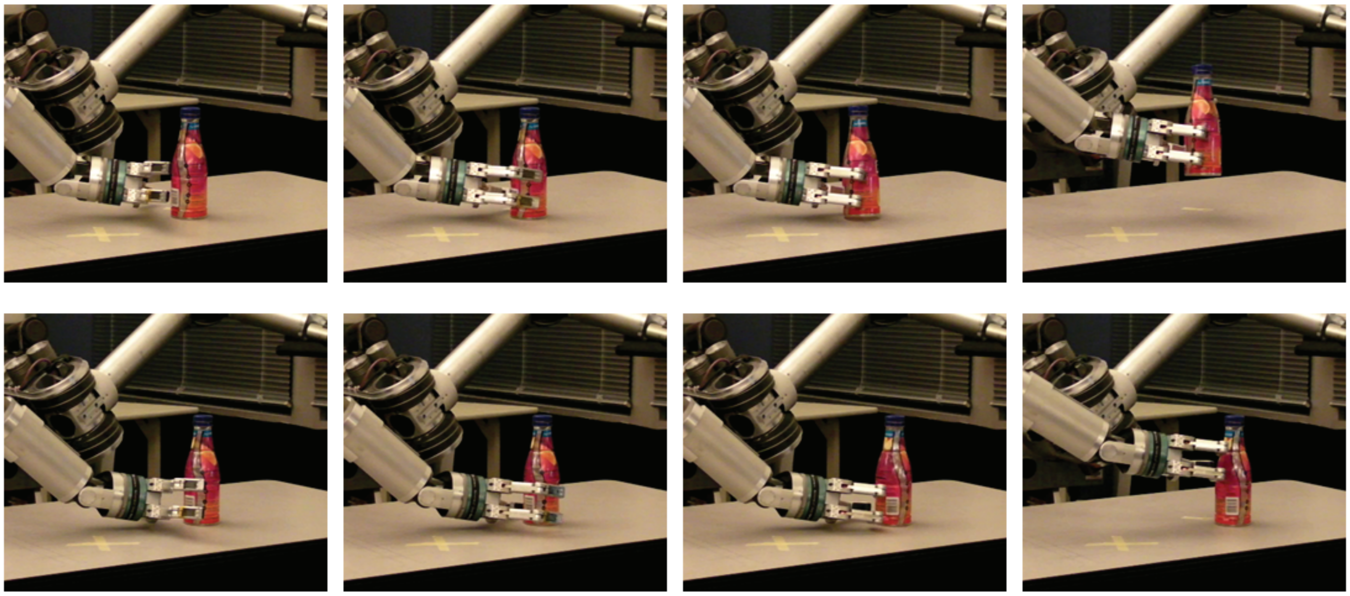

Fig. 2. Grasp success (upper) and failure (lower) cases with providing the same grasp to the planner. The two actual hand poses (relative to the object) before closing fingers differ from the planned pose due to the data uncertainty.

pose and contact points during grasping, by which we expect to be able to evaluate the grasp quality more precisely. In addition, at every grasp test, we run multiple simulations starting from slightly different initial conditions sampled from a probability model to consider the uncertainty, and then evaluate the success rate of the grasp by gathering all simulation results. As far as we know, this is the first attempt to incorporate both dynamics and data uncertainty in estimating the grasp success rate.

The key finding of our work is that both the dynamic contact change at the moment of grasping and the object pose error from the sensor uncertainty must be considered at the same time in grasp quality evaluation. Our new method considering both of them outperformed the existing method in predicting the actual (experimental) grasp success rate, while considering only one factor at a time did not improve performance.

II. RELATED WORK

Much previous research related to grasp quality metrics has been focused on analyzing the 6-dimensional space spanned by contact wrenches. Li and Sastry [1] proposed using the smallest singular value of the wrench space matrix and the volume of the wrench space as quality metrics. They also suggested a task oriented quality measure to consider the type of task to be done with the grasp. Ferrari and Canny [2] suggested to use the radius of largest ball inscribing the convex hull of the contact wrenches. The physical meaning of this metric is that it represents the largest minimum disturbance wrench that can be resisted by the contacts. This 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! [3]. A variety of other grasp quality metrics have been considered, and are summarized nicely in [17]. In this article, Balasubramanian et al. additionally consider grasp quality measures that may be derived from human-guided robot grasps.

Grasp quality measures have many applications such as finding an optimal grasp as illustrated in [1], [2]. Automatic generation of a large grasp set is also an obvious application. A large grasp database containing grasps of various objects was built using GraspIt! [5]. Other manipulation planning tools such as OpenRAVE [4] also provide a function to generate a grasp set automatically because the precomputed grasp set can be used in motion planning algorithms such as RRT-Connect [6]. Most of the existing methods for automatic generation of a grasp set first obtain the contact points of the grasp through a simple kinematic simulation with an assumption that the object is static, i.e., that is remains at the same position even after collision with the fingers. However such a static object assumption does not hold in many practical situations. As shown in Fig. 2, the object can move significantly in response to collision with the fingers, and in such cases the existing metrics may not give us useful information on grasp quality. To handle the issue, we consider dynamics of the 3-dimensional object in our simulation for evaluating grasp quality.

Handling data uncertainty in grasping has been studied by many researchers. A number of researchers (e.g.,[14], [13]) have considered planning actions in conditions of complete uncertainty. However, we assume that our robot is capable of estimating object configuration with moderate error. Toward handling such a circumstance, Zhen and Qian consider how small uncertainties in friction coefficient and contact locations affect grasp quality. Christopoulos and Schrater [7] 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. [15] address shape uncertainty by cross testing grasps with alternative shapes that are nearest neighbors to a given model. It is worth noting that their simulator uses an approximate dynamics simulator to compute object response (e.g., object motion due to forces applied by the hand is captured, but other environment forces such as the supporting surface are not modeled). The simulator used in their research has been incorporated into GraspIt! as described in [16]. Hsiao et al. [9] introduced a method that takes into account uncertainty in object shape and pose data. They combined the data from a set of object detection algorithms using a probabilistic framework to find an optimal grasp.

In this paper we consider data uncertainty such as object pose, with a focus on capturing the effect of errors in pose estimation on the result of dynamics simulation. To our knowledge, this is the first work to address both pose uncertainty and the dynamic effect of the grasping task.

III. GRASP QUALITY EVALUATION

We define a grasp as a combination of a relative pose of the hand to the object and the initial finger joint configuration. The grasp information is used to set the initial condition of the hand in our physically based simulation.

In each simulation we apply a simple grasping motion to the hand to close the fingers for object grasping and to lift up the object while maintaining the closed finger configuration, and compute the object motion using the equations of motion of the rigid body. In order to evaluate the quality of a single grasp, we run multiple simulations, and at each run, we set a slightly different object initial pose obtained by sampling from a probability distribution model to consider the uncertainty in the sensor based pose localization. At each simulation we measure the quality of the simulated grasp and gather the quality scores of all simulations we ran for a particular grasp as raw data for evaluation of the quality of the grasp under uncertainty.

One way to measure the quality of a grasp from the simulation results would be simply counting the number of successful grasps from the full set of simulation variations for that grasp. 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 the lift-up phase, and give a score of 0. If the object did not move and it had contacts with more than three links after the lift-up, we regard the grasp to be successful and give a score of 1. If the object was still moving after completion of the lift-up phase or it had contacts with only two links, then we give a score of 0.5. We used the same scoring system to measure the quality, or success rate, of a grasp in the real robot experiment too. See section V for comparison of our simulation results with those of the experiment.

We also tested another quality measure that considers movement of the object as it is being lifted from its original supporting surface. We mark the object pose relative to the hand as a reference after completion of the grasping phase and before starting the lift-up. Then, we obtain the maximum

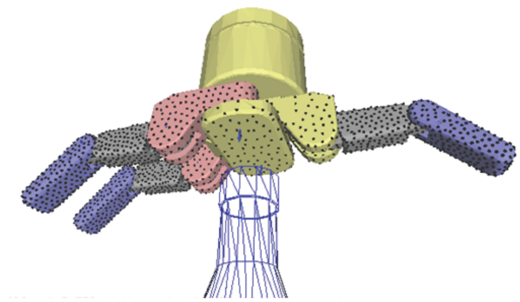

Fig. 3. Uniformly distributed surface points for collision detection

deviation of the relative object pose from the reference pose during the lift-up. If the object was moving a lot during the lift-up, we consider the grasp is not stable and give it a low score value. On the contrary, if there was no relative object motion, we give a score of 1 which is the highest in our system. Of course, as in the case of the above measure, if the object was out of the hand or it had contacts with less than two links at the final time step, then we still regard it as a failure case and give a score of 0 which is the lowest. We consider the deviations in the position and orientation separately and compute them with

$$\delta_p = ||p_{com} - \bar{p}_{com}|| \tag{1}$$

$$\delta_R = ||Log(\bar{R}^T R)|| \tag{2}$$

where p_{com} and R denote the relative center of mass position and the orientation of the object from the hand, and the bar symbol represents the reference values. Then, the quality score is obtained from the maximum deviation $\delta^{\max}$ and a tolerance limit L as

$$q = 1 - \frac{\delta^{\max}}{L} \tag{3}$$

when $\delta^{\max} < L$. If the maximum deviation exceeds the tolerance limit, we give a score of 0. This quality metric can be computed separately for position and orientation and used as such, or these metrics can be combined, for example by averaging. In our simulation we set 1 cm and 10 deg as the tolerance limits, and the effect of change in the limits on the relative grasp quality score was negligible.

IV. IMPLEMENTATION DETAIL

We simulate the robot hand grasping an object using our own physics simulation tool. We assume frictional contacts between the object and hand, and between the object and the ground. The object motion at every time step is computed based on the current contact forces and the equations of motion of the rigid body.

The finger joints are kinematically driven by motors to close the fingers, and the motor speed is adjusted depending on the magnitude of the motor torque. The motor torque is obtained by converting the contact forces using Jacobian matrices, and hand dynamics is not considered in our simulation. As the motor torque increases due to forces at

Fig. 4. The HERB robot and the objects (Pop Tarts and Fuze Bottle) used in the test. The objects are fully filled with contents.

the contacts, the motor speed is reduced linearly which is a typical characteristic of DC motors. If the motor torque exceeds a limit, the motor does not close the finger any more and remains in its current position. If all fingers have been closed, then the hand starts to lift up the object by following a given trajectory and stops at some point. We use a trapezoidal velocity profile to lift the hand.

Some robotic hands have a clutch mechanism to change the finger closing behavior, and our system handles such a clutching mechanism. In the case of Barrett hand which was used in our simulation and experiment, a single motor usually drives the two joints in each finger. But, after the breakaway clutch is triggered, only the outer joint rotates, closing the finger tip further while the other joint remains in place, which is useful for closing the hand more tightly around an object.

For collision detection, we use a set of points uniformly distributed on the surface of the robot hand geometry (Fig. 3). Each point will be tested for collision with the surface of the object, and at each contact point, a contact force will be computed with a penalty-based contact model. We use a penalty-based frictional contact model suggested by Yamane and Nakamura [10]. The contact forces affect the state of the object at the next time step through the equations of motions of the rigid body. They also affect the motor speed for closing the fingers so that the motors automatically adjust their speed and eventually stop at the stall torque condition.

V. RESULTS

We tested our grasp quality measure with predefined grasps for the two types of objects shown in Fig. 4. We generated a very large set of grasps using OpenRAVE and chose 10 grasps for each object as shown in Fig. 5. The grasp quality of the grasps were estimated using our simulation based algorithm described in Section III.

To consider the uncertainty in the object pose relative to the hand, we sampled the position and orientation of the

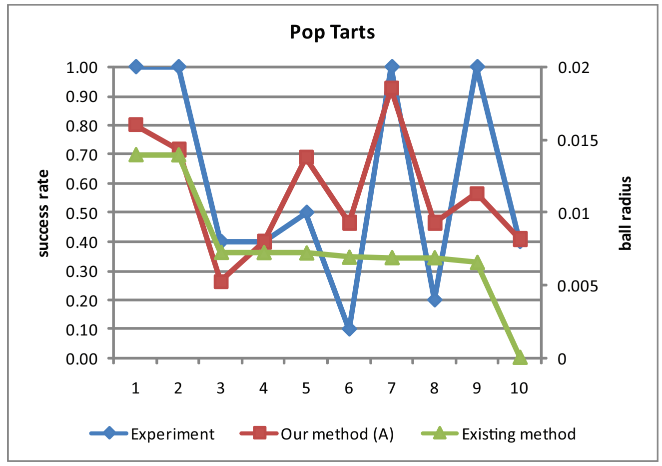

Fig. 6. Comparison of the quality measures (Pop Tarts)

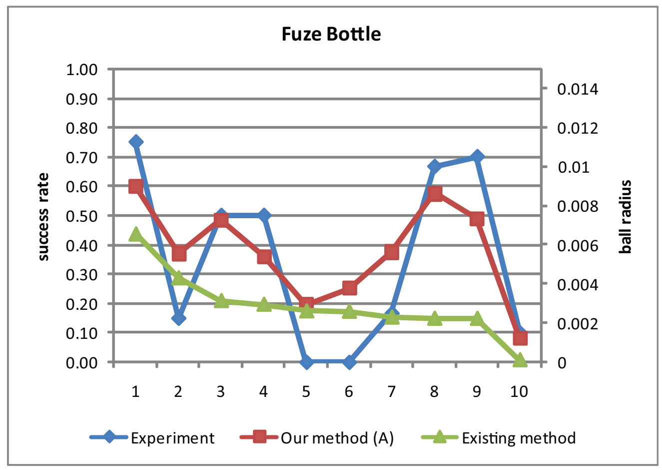

Fig. 7. Comparison of the quality measures (Fuze Bottle)

object using normal distribution models. We set the standard deviations of the two distribution models to be 1 cm and 6 deg respectively, and the sample size was 100 for each grasp in the test. It took about 5 min and 13 min to evaluate the 100 trials for each of 10 grasps (total of 1000 simulations per each object) for the box shaped object (Pop Tarts) and the bottle (Fuze Bottle) respectively, and the test was performed on a laptop with a 2.67 GHz Intel Core i7 CPU.

We conducted an experiment to obtain actual success rates, or quality, of the grasps with HERB (Fig. 4), a service robot consisting of two Barrett WAM arms and Barrett hands [11]. For each grasp, we repeated 10 times to grab and lift up the objects to measure the success rate of the grasp, marked a score at each trial using the (0, 0.5, 1) scoring system described in Section III, and averaged the scores to get an average success rate of the grasp.

Fig. 6 and 7 shows the grasp quality values for the two objects. The blue line with diamond marks shows the success rate obtained by the experiment, the red line with squares represents the grasp quality values measured by our method using the (0, 0.5, 1) scoring system, and the green line

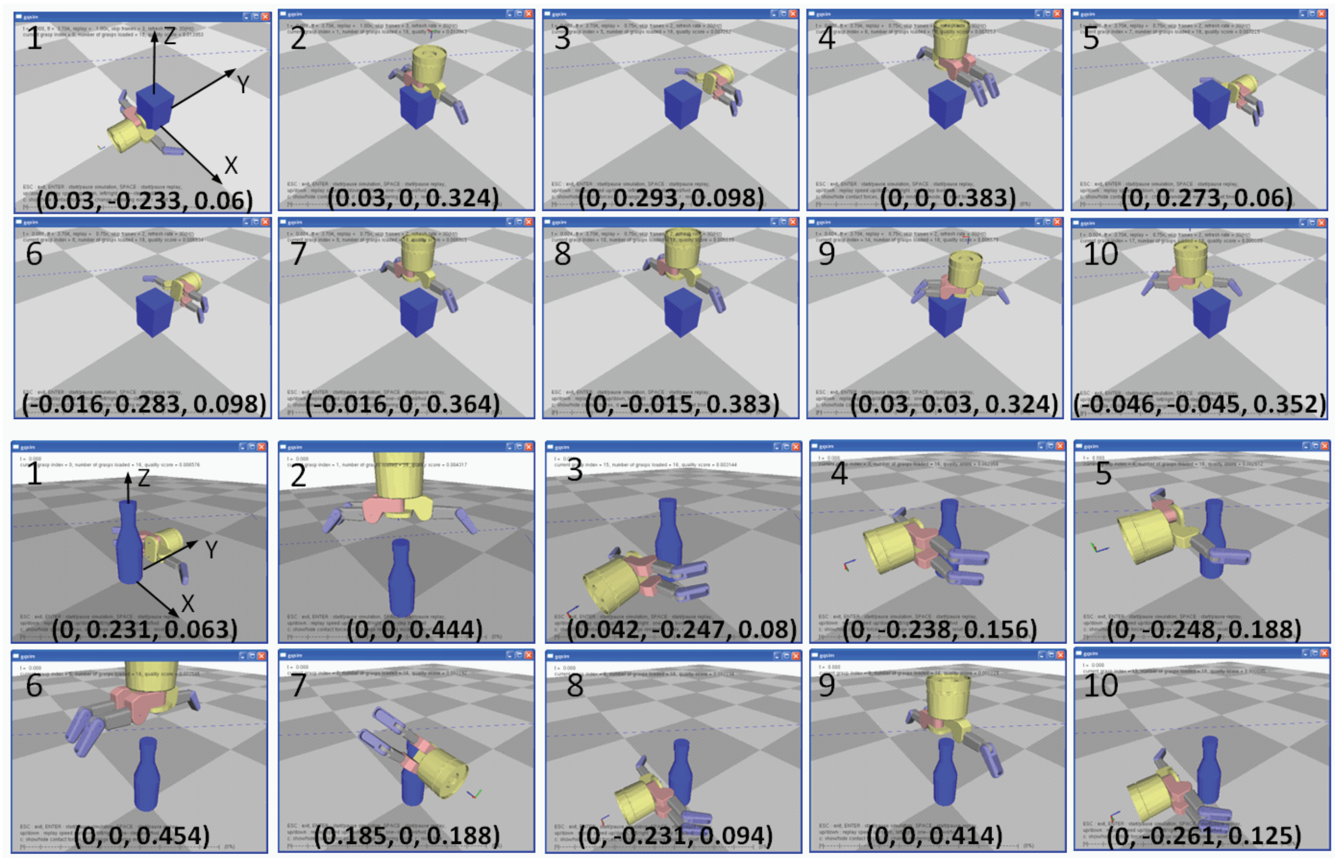

Fig. 5. The grasps tested in our simulation and experiment. Relative position of the hand to the object is shown at the bottom of each grasp.

with triangles shows the existing grasp quality measure by Ferrari and Canny [2] with the static object assumption. Though the quantity of the values is somewhat different, our grasp measure shows a strong tendency of following the experiment results, and we can conclude that the method can be used effectively in predicting the actual grasp success rates.

In our experiment, the actual robot grasping tends to fail easily especially when the grasp is located at the near boundary of the area where the finger can reach the object. For example, the grasps 6 and 8 of the Pop Tarts, and the grasps 2 and 6 of the Fuze Bottle are located far from the object, and the success rates in the actual robot grasping are significantly lower than the simulated values. We think that, in such cases, the success rate in the experiment is more likely to be affected by a calibration error. If the pose error distribution has a nonzero mean, and the offset is located in the far side, then the robot would have only small chance of successful grasping. Of course, the offset could be placed in favorable side with an equal probability if the object is placed randomly. However, we used only the right arm of the robot in the experiment, so this restricted the operator from placing the object fully randomly because of the limited work space of the robot arm. On the contrary, we assumed that the mean of the sensory error distribution is zero in our simulation, so the simulated grasping would still have 50% of chance for having the hand be located in the near side of the object, and so we speculate that this leads to a higher success rate than the actual grasping.

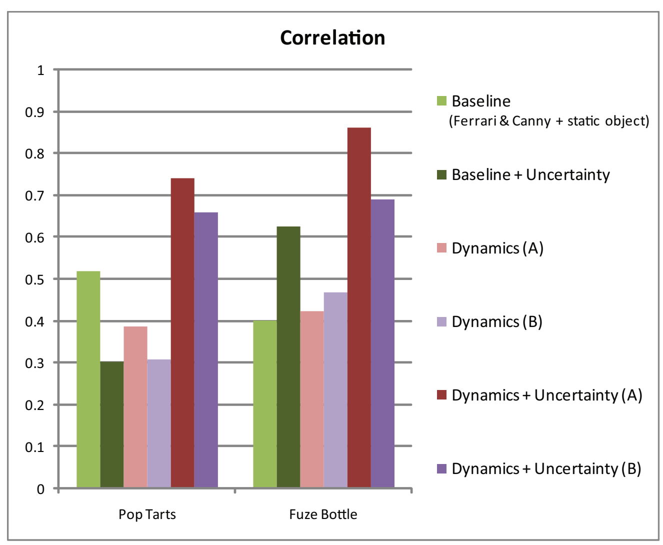

Fig. 8. Correlations to the experiment data

In Fig. 8 we show the statistical correlations between the experiment data, i.e., the actual grasp success rate, and a few grasp quality measures. In addition to the two grasp measures discussed above (marked as 'Dynamics + Uncertainty (A)' and 'Baseline'), we added our another version of quality measure based on the maximum deviation (marked as 'Dynamics + Uncertainty (B)'). In our experiment, our first measure using the three-step scoring system showed better performance than the second method in predicting the actual grasp success rates. We expect that the first measure performs better because we used the same scoring system to measure the success rates in the experiment.

We also added the quality scores of our two measures without considering the uncertainty model (marked as 'Dynamics (A)' and 'Dynamics (B)') — we execute a single run of physics simulation per grasp and get the scores as described in Section III. The correlation plot clearly shows the value of adding pose uncertainty in the grasp quality measure design. The two measures incorporated with the uncertainty model show the highest correlation to the actual grasp success rates, which shows the effectiveness of our grasp measures in estimating the actual grasp quality. However, the quality scores without pose uncertainty do not give good information on the success or failure of the actual grasp even with considering object dynamics. From the observation, we speculate that modifying the existing method by simply replacing the static object assumption with a dynamics simulation would not produce a great improvement in estimating the real grasp success rates.

Finally, we investigated the effect of considering object pose uncertainty on the performance of the existing quality measure by testing poses sampled from the uncertainty model and averaging the scores for each grasp (marked as 'Baseline + Uncertainty'). However, this strategy did not produce consistent improvement in our examples — it raised the correlation in the case of the Fuze Bottle but acted in the opposite way for the other object.

VI. CONCLUSION

In this paper two core factors affecting success or failure in actual robotic grasping — contact change at the moment of grasping and the object pose error by the sensor uncertainty – have been addressed, and new grasp quality measures based on physics simulation have been suggested. The effect of changing contacts is captured by adding object dynamics in the simulation, and the object pose uncertainty is considered by running multiple simulations with slightly different initial poses obtained from a probabilistic distribution model. Through the real robot experiment, we showed that the new methods can effectively estimate the actual grasp success rates. We also have found that both the dynamics and uncertainty must be considered at the same time in grasp quality evaluation, and considering only one factor at a time does not improve the performance of predicting actual grasp success rates.

Our new methods show better performance in predicting the actual grasp success rates than the existing method, and this would help improving robustness of robot grasping by providing a better grasp set to the planner. However, the methods based on either counting the number of finger contacts or measuring the object pose deviation during the lift-up process do not explicitly tell us about the stability of the final grasps, i.e., the ability to resist external disturbance forces, which is important for achieving robust manipulation. It would be an interesting future work to design a simulation-based method for measuring grasp stability and test it by comparing with experimental results.

REFERENCES

[1] Li and Sastry, "Task-oriented optimal grasping by multifingered robotic hands," *IEEE Journal of Robotics and Automation*, 1988.

[2] Ferrari and Canny, "Planning optimal grasps," *IEEE International Conference on Robotics and Automation*, 1992.

[3] Miller and Allen, "Graspit! A versatile simulator for robotic grasping," *IEEE Robotics and Automation Magazine*, 2004.

[4] Rosen Diankov, "Automated Construction of Robotic Manipulation Programs," doctoral dissertation, tech. report CMU-RI-TR-10-29, Robotics Institute, Carnegie Mellon University, September, 2010.

[5] Corey Goldfeder, Matei Ciocarlie, Hao Dang, Peter K. Allen, "The Columbia Grasp Database," *IEEE International Conference on Robotics and Automation*, 2009.

[6] James J. Kuffner Jr. and Steven M. Lavalle, "RRT-Connect: An efficient approach to single-query path planning," *IEEE International Conference on Robotics and Automation*, 2000.

[7] Vassilios N. Christopoulos and Paul Schrater, "Handling shape and contact location uncertainty in grasping two-dimensional planar objects," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.

[8] Kaijen Hsiao, Tomas Lozano-Perez and Leslie Pack Kaelbling, "Relatively Robust Grasping," *ICAPS 2009 Workshop on Bridging the Gap Between Task and Motion Planning*, 2009.

[9] Kaijen Hsiao, Matei Ciocarlie and Peter Brook, "Bayesian Grasp Planning," *ICRA 2011 Workshop on Mobile Manipulation: Integrating Perception and Manipulation*, 2011.

[10] Katsu Yamane and Yoshihiko Nakamura, "Stable penalty-based model of frictional contacts," *IEEE International Conference on Robotics and Automation*, 2006.

[11] S. Srinivasa, D. Ferguson, C. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M. VandeWeghe, "HERB: A Home Exploring Robotic Butler," *Autonomous Robots*, 2009.

[12] Yu Zheng and Wen-Han Qian, "Coping with the Grasping Uncertainties in Force-closure Analysis," *International Journal of Robotics Research*, 24(4), 2005.

[13] K. Y. Goldberg and M. T. Mason, "Bayesian Grasping," *IEEE International Conference on Robotics and Automation* 1990.

[14] R. C. Brost, "Planning Robot Grasping Motions in the Presence of Uncertainty," *The Robotic Institute, Carnegie-Mellon Tech. Report CMU-RI-TR-85-12*, 1985.

[15] C. Goldfeder and P. K. Allen, "Data-driven grasping," *Autonomous Robots*, 31(1), 2011.

[16] 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), 2005.

[17] R. Balasubramanian, L. Xu, P. D. Brook, J. R. Smith, and Y. Matsuoka, "Human-Guided Grasp Measures Improve Grasp Robustness on Physical Robot," *IEEE International Conference on Robotics and Automation*, 2010.