Abort and Retry in Grasping

Alberto Rodriguez, Matthew T. Mason, Siddhartha S. Srinivasa, Matthew Bernstein, and Alex Zirbel
The Robotics Institute — Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213, USA
{albertor, matt.mason, siddh}@cs.cmu.edu
{mbernste, azirbel}@andrew.cmu.edu

***Abstract*— Iteration is often sufficient for a simple hand to accomplish complex tasks, at the cost of an increase in the expected time to completion. In this paper, we minimize that overhead time by allowing a simple hand to abort early and retry as soon as it realizes that the task is likely to fail. We present two key contributions. First, we learn a probabilistic model of the relationship between the likelihood of success of a grasp and its grasp signature—the trace of the state of the hand along the entire grasp motion. Second, we model the iterative process of early abort and retry as a Markov chain and optimize the expected time to completion of the grasping task by effectively thresholding the likelihood of success. Experiments with our simple hand prototype tasked with grasping and singulating parts from a bin show that early abort and retry significantly increases efficiency.**

I. Introduction

Simple hands are characterized by simple mechanical designs and simple control strategies, both of which compromise the potential generality of the hand. In practice, and based on observations of humans using simple tools and effectors, simple hands offer broader manipulation capabilities than any autonomous system has yet demonstrated.

After arguing the case for simplicity in [1], and with the aim of demonstrating manipulation capabilities with simple hands, we approached the problem of singulating objects out of a bin in [2]. The approach in [2] has three key elements:

- **Simple mechanism**: The simple hand in Fig. 1 has thin cylindrical fingers compliantly coupled to a single actuator, arranged symmetrically around a low friction flat circular palm.
- **Simple control**: Contrary to the more traditional approach where complex robot hands that try to "put the fingers in the right place", we close the hand and "let the fingers fall where they may.". We expect the fingers either to drive the object to a stable pose or to reject it, effectively reducing the space of possible outcomes of the grasp. By simplifying the relationship between the signature of a grasp and its outcome, we facilitate the creation of a data-driven model.

This work was sponsored in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-10-2-0016, DARPA-BAA-10-28, NSF-IIS-0916557, and NSF-EEC-0540865. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, DARPA, or the U.S. Government.

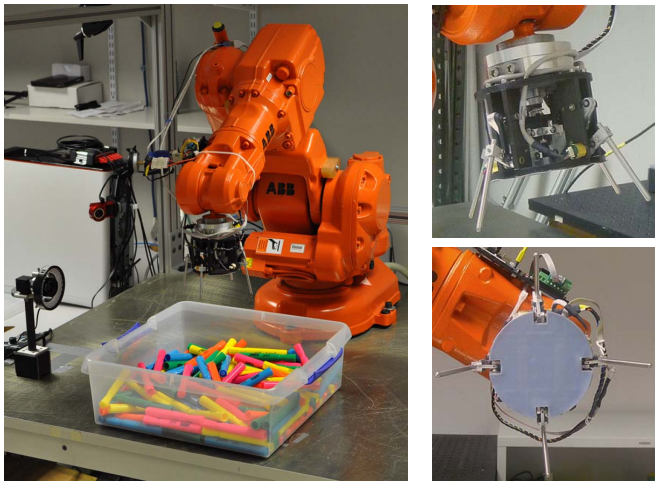

(a) (b)

Fig. 1. (a) Bin-picking scenario. A robotic manipulator drives a simple hand in and out of a bin full of identical objects. (b) P2, a simple hand prototype with four fingers compliantly coupled to a single actuator. Motor and fingers have absolute encoders that provide the full state of the hand.

- **Iteration**: To address the expected shortcomings of the simple approach, we iterate upon failure (Fig. 2) until we succeed in singulating a part.

The simplicity of the approach often comes at the cost of increasing the expected time to a successful grasp. The robot reaches into the bin of parts (Fig. 1), closes its hand and pulls it out, after which it queries its encoders and uses a learned discriminative model to decide if it has a single part. If it does not, it iterates until it succeeds.

We propose to reduce the expected time to completion by using proprioceptive feedback to predict failure early during execution and possibly abort and restart the procedure. Key to our approach is the concept of *grasp signature*, the trace of the state of the hand along the entire grasp motion (Sect. III). We learn a probabilistic model to track the instantaneous probability of success during the grasp process. Based on that probability, we model abort and retry as a Markov chain, and derive analytically the expected time to a successful grasp (Sect. IV). We then use the model to optimize abort thresholds to minimize the expected time to a successful grasp (Sect. V).

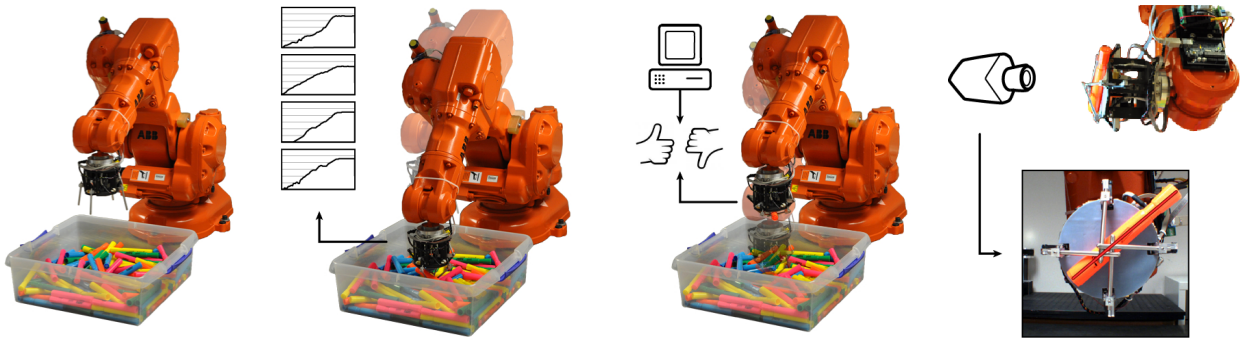

Fig. 2. Iterative framework: The grasp is iterated until the system predicts a high probability of success. (a) At the beginning of the grasp the hand is above the bin. (b) The interaction of the hand with the environment is logged during the entire grasp process, constituting the grasp signature. (c) As the grasp is being executed, a data-driven model classifies the grasp as success or failure. In the case of failure, the grasp immediately aborts and restarts. (d) During the training phase, the learning process is supervised by a vision system that provides ground truth on the number and pose of grasped markers.

Our experiments show that early abort and retry significantly increases the efficiency of the system (Sect. VI). Furthermore, we demonstrate how the overhead of the iterative framework—expected time to a successful grasp minus the nominal length of the grasp—evolves with the number of possible abort points along the grasp process (Sect. VII).

II. Previous Work

We focus on a bin-picking scenario, a challenging grasping task that combines high clutter with high pose uncertainty, and for decades the focus of numerous research efforts [3], [4].

We use a data-driven approach for failure detection and abort optimization based on the signature of the grasp. Data-driven approaches have previously been proposed for failure detection in different contexts, including milling operations [5], [6], machine vibration analysis [7] or failure detection in automated assembly [8], [9].

Dollar et al. [10] showed how in-hand sensor information can be used to improve grasping performance. In the process of detecting failure, we estimate the outcome of the grasp based on kinesthetic sensor data. Bicchi, Salisbury and Brock [11] explored a similar problem: assuming known finger shape and location, they estimate the contact point from a measured applied wrench, a technique known as *intrinsic contact sensing*. This contact information can be used to infer the pose of a known shape.

Also relevant is the related problem of inferring object location from kinesthetic or contact data, studied by Siegel [12], Jia and Erdmann [13], [14] and Wallack and Canny [15]. All of these works are based on analytical models of contact and grasp mechanics. Instead, we use a statistical data-driven approach to create a model of the complex relationship between the signature of the grasp process and its outcome, thereby bypassing the intermediate estimation of contact points. With this, we are able to incorporate numerous sources of information that are very challenging to model, including the effect of the grasping motion and that of surrounding clutter. Laaksonen, Kyrki and Kragic [16] compared the effectiveness of different statistical data-driven methods for estimation of grasp stability, based on both kinesthetic and contact sensor data.

Our framework for optimizing the series of abort conditions has some similarities with the cascades of classifiers proposed in computer vision to speed up the detection rate of object classifiers without compromising performance [17], [18]. A series of incrementally more computationally expensive classifiers trade off the cost and risk of taking a decision or letting the following classifier in the cascade do it.

III. Preliminary Concepts

A. Grasp Signature

We define the *grasp signature* $\mathcal{G}$ as the trace of the state of the hand along the entire grasp motion as perceived by the hand's sensors. The signature can be composed of, but not limited to, time-stamped data from joint encoders, tactile sensors, and torque sensors. This work tests the hypothesis that the grasp signature encodes enough information to characterize the outcome of a grasp.

The compliant simple hand we have built, called P2, has absolute motor and fingers encoders that allow us to recover the full kinematic state of the hand. Figure 3 shows a side by side comparison of the finger encoder signatures of examples of successful and failed grasps.

B. Expected Time to a Successful Grasp

We define the *expected time to a successful grasp* τ, as the expected number of attempts until a successful grasp multiplied by the time span of the grasp T.

In [2] we improve system precision (false negative rate) by tuning the weights of positive and negative training examples, at the cost of increasing the false positive rate. This has the unsought consequence of increasing the *probability of iteration* f of the system, leading to an increase in the expected time to a successful grasp.

In this paper we propose to abort grasps that the system predicts likely to fail as a technique to reduce the expected

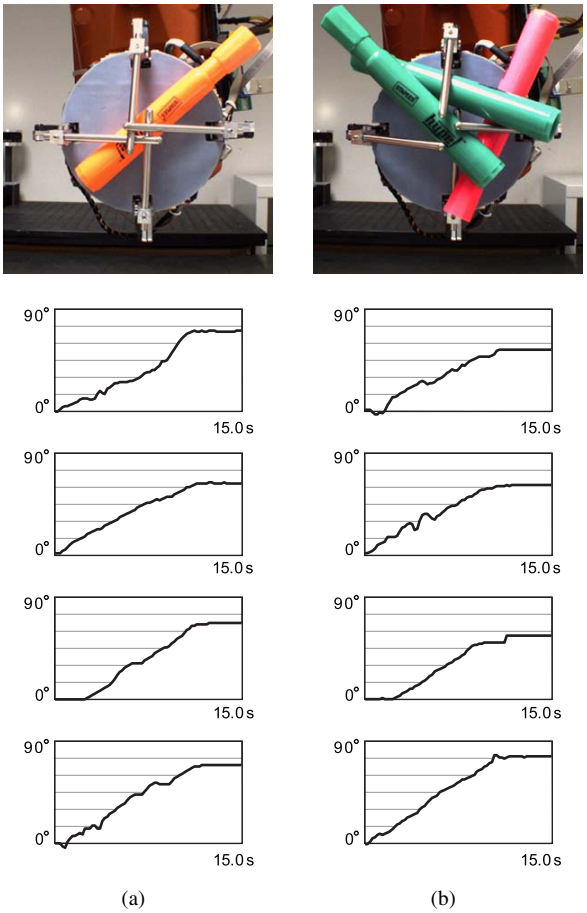

Fig. 3. Side by side comparison of the grasp signature (4 finger encoders) of a typical (a) successful and (b) failed grasp. The fingers begin the grasp perpendicular to the palm (0°) and reach the final position shown in the figures.

time to a successful grasp while still maintaining high values for the precision of the system.

The expected time to a successful grasp of the simple system in Fig. 2 is related to the iteration probability by $\tau = \frac{T}{f}$. Analogously we show in Sect. IV that, for the more the complex system that allows early abort, the expected time is a function of the probability of early failure.

C. Success Probability

Our goal is to decide at time t whether the grasp is worth continuing, or if we should abort. To do so, we learn a probabilistic model of the relationship between the first t seconds of the signature and the expected final outcome of the grasp. This contrasts with our previous approach [2] where we learn a discriminative model to signal success only at the end of the grasp.

We take as input a set of K grasp signatures $\{\mathcal{G}_i(t)\}_{1\ldots K}$ and correspondent labels of success or failure and learn a probabilistic model of their relationship:

$$\mathcal{M} : \mathcal{G}([0,t]) \mapsto p(t)$$

where $\mathcal{G}(t)$ is the signature of a grasp execution and $p(t)$ is the instantaneous probability of success of the grasp as predicted by the model $\mathcal{M}$. As an example, Fig. 4 shows the success probability signal for the successful and failed grasp signatures in Fig. 3. At the beginning of the grasp, the system is uncertain about the possible outcome of either grasp, but becomes more confident over time. We describe the model learning algorithm in Sec. VI-C.

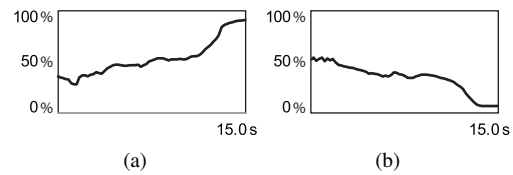

Fig. 4. Evolution of the probability of success for the successful (a) and failed (b) examples illustrated in Fig. 3.

It is impractical to track the probability of success continuously. In practice we discretize the grasp into n time slices $[t_0, t_1], [t_1, t_2] \ldots [t_{n-1}, t_n]$ and train n independent probabilistic classifiers at instants $\{t_i\}_{1\ldots n}$. As the grasp progresses, they output a series of estimated success probabilities $p_1 \ldots p_n$. To make a decision to abort, we compare these against n cutoff probability thresholds $\pi_1 \ldots \pi_n$:

$$\text{If } p_i < \pi_i \;\rightarrow\; \text{ABORT at } t_i \tag{1}$$

A unique contribution of our paper is an analytical model of the execution of such a system with n possible aborting points (Sect. IV), and an optimization technique for tuning the cut-off probabilities to minimize the expected time to a successful grasp based on training data (Sect. V).

IV. Modeling Abort and Retry

In this section we model the steady-state behavior of the system with n abort points. We call state O to the beginning of the grasp motion, states $S_1 \ldots S_n$ to the abort points at instants $t_1 \ldots t_n$, and state S to the final successful grasp.

We represent the system by the graphical model in Fig. 5. The system satisfies the Markov property, i.e., future state depends only on the current state and not on the past. All states except the initial one trivially satisfy the property since, beginning from O, there is only one possible series of transitions to get to them. If we assume statistical independence between successive repetitions of the experiment, the initial state also satisfies the Markov property.

The behavior of a time-homogeneous Markov chain is represented by the transition probabilities between states, which we will call: probability of transition from state S_i to state S_{i+1}. These are related to (1) by:

$$P_i = P\,[\; p_i > \pi_i \;] \tag{2}$$

At each state S_i the abort probability is then $1 - P_i$. P_n is the fraction of grasp attempts that reach the end of the execution and are actually classified as good grasps.

The Markov chain model proves useful for analyzing the steady-state behavior of the system and, in particular, for computing the expected time to a successful grasp. We

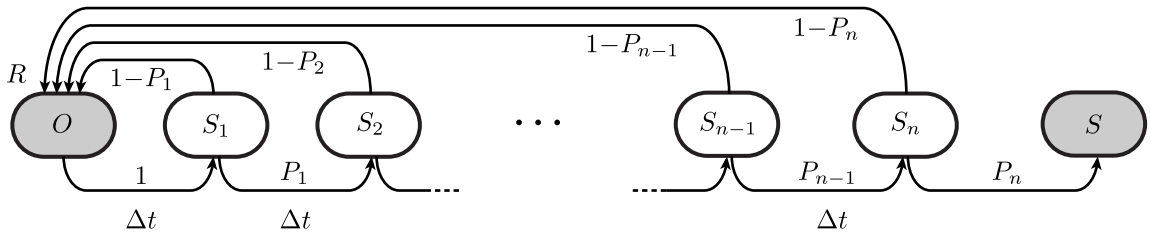

Fig. 5. Markov Chain that models a system with n abort points at instants $t_i = i \cdot \Delta t$. State O represents the beginning of the grasp, states $S_1 \dots S_n$ are the n abort points and S the final successful grasp. The system reaches the end state if and only if it is not discarded by any of the classifiers at the abort points. The cost of forward transition is equal to Δt and we assume that the cost of aborting a grasp is independent of the state and equal to R. At each state S_i the grasp continues with probability P_i and aborts with probability $1 - P_i$.

assume all abort points to be equispaced in time, with a constant spacing of Δt. We also suppose that the time cost of aborting, R, is constant and independent of the state of the system. The following proposition gives a closed form expression for the expected time to success as a function of the transition probabilities.

Proposition 1 (Expected time to a successful grasp): In the iterative system of Fig. 5 with n equispaced abort points and transition probabilities $P_1 \dots P_n$, the expected time to success τ can be expressed as:

$$\tau = \Delta t \left[\frac{1 + \sum_{i=1}^{n-1} \left(\prod_{j=1}^{i} P_j \right)}{\prod_{i=1}^{n} P_i} \right] + R \left[\frac{1 - \prod_{i=1}^{n} P_i}{\prod_{i=1}^{n} P_i} \right] \quad (3)$$

Proof: We introduce the intermediate variables τ_0, $\tau_1 \dots \tau_n$ to represent the expected time to success from each one of the states of the system O, $S_1 \dots S_n$ correspondingly. Notice that τ_0 is by definition equal to the expected time to a successful grasp τ. We will prove the general term by induction on the number of time slices n.

For the case $n = 1$, the Markov chain reduces to the one in Fig. 6. This corresponds to the original framework in Fig. 2 where the decision to abort is taken solely at the end of the grasp.

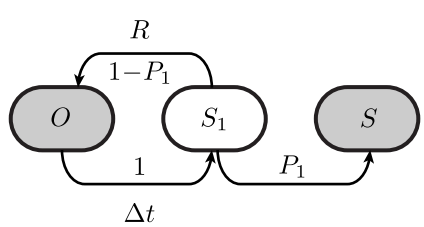

Fig. 6. Markov chain of the system for the case of $n = 1$. In this case, the only possible abort point t_1 is at the end of the grasp.

In steady-state the expected times τ_0 and τ_1 are related by the equations:

$$\begin{cases} \tau_0 &= \Delta t + \tau_1 \\ \tau_1 &= P_1 \cdot 0 \; + \; (1 - P_1) \cdot (R + \tau_0) \end{cases} \quad (4)$$

Solving the system (4) for τ_0, we get:

$$\tau_0 = \Delta t \cdot \frac{1}{P_1} + R \cdot \frac{1 - P_1}{P_1} \quad (5)$$

which satisfies the general term.

Now we assume that the general term is correct for the case of $n - 1$ abort points and we prove for the case of n. The Markov chain with n abort points in Fig. 5 is equivalent to the simplified chain in Fig. 7 where the first $n - 1$ states are combined into a macro initial state O^* with transition cost τ^*.

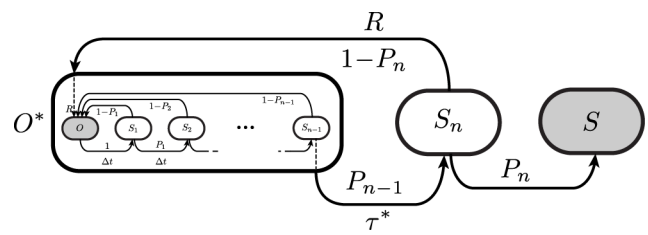

Fig. 7. Equivalent Markov chain for the case of n abort points. The first $n - 1$ abort states can be included in a macrostate O^*. The new transition cost τ^* from O^* to S_n is the expected time to a successful grasp of the subsystem with $n - 1$ abort states plus Δt, the original transition cost from S_{n-1} to S_n.

The new initial state behaves internally as a system with $n - 1$ abort points. By induction, the transition cost from O^* to S_n is:

$$\tau^* = \Delta t \left[\frac{1 + \sum_{i=1}^{n-2} \left(\prod_{j=1}^{i} P_j \right)}{\prod_{i=1}^{n-1} P_i} \right] + R \left[\frac{1 - \prod_{i=1}^{n-1} P_i}{\prod_{i=1}^{n-1} P_i} \right] + \Delta t$$

The simplified equivalent system in Fig. 7 has the same structure as in the case of $n = 1$, therefore, using (5) the expected time to a successful grasp can be computed as:

$$\begin{aligned} \tau_0 &= \tau^* \cdot \frac{1}{P_n} + R \cdot \frac{1 - P_n}{P_n} = \\ &= \Delta t \left[\frac{1 + \sum_{i=1}^{n-2} \left(\prod_{j=1}^{i} P_j \right)}{\prod_{i=1}^{n} P_i} \right] + \\ &\quad + R \left[\frac{1 - \prod_{i=1}^{n-1} P_i}{\prod_{i=1}^{n} P_i} \right] + \Delta t \cdot \frac{1}{P_n} + R \cdot \frac{1 - P_n}{P_n} = \\ &= \Delta t \left[\frac{1 + \sum_{i=1}^{n-1} \left(\prod_{j=1}^{i} P_j \right)}{\prod_{i=1}^{n} P_i} \right] + R \left[\frac{1 - \prod_{i=1}^{n} P_i}{\prod_{i=1}^{n} P_i} \right] \end{aligned}$$

which concludes the proof. ■

In the next section we see how the expression for τ in Proposition 1 simplifies the estimation of the expected time to a successful grasp of the system.

V. Optimizing Abort and Retry

As detailed in Sect. III-C, the learning system comprises n predictive models producing success probability estimates p_i at abort points t_i, and abort thresholds π_i. In this section we show how to optimize those thresholds to minimize the expected time to a successful grasp τ.

We first need to study how variations of the thresholds π_i map to changes in the expected time τ. An online experimental approach would be impossibly time consuming, requiring numerous experiments to estimate the expected time to success for every queried value of the thresholds.

Instead of the experimental approach, we combine the analytical expression of Proposition 1 with offline experiments. Given a candidate set of thresholds π_i, we can use offline experimental data to estimate transition probabilities P_i, and then apply equation (3) to estimate τ. The transition probabilities can be estimated experimentally by running K grasp executions and computing:

$$P_i = P\left[\ p_i > \pi_i\ \right] = \frac{\text{Grasps reach } S_{i+1}}{\text{Grasps reach } S_i} \tag{6}$$

It is key to notice that, when using the transition probabilities as an intermediate step to evaluate τ, it is not necessary to run the experiment again when the values of the cut-off probabilities change. Equation (3) allows for a more efficient strategy. Assuming that the learned predictive models of the probability of success do not change, we can reapply the abort condition in (1) to the same set of grasp executions, but now with the new cut-off probabilities. We then reestimate the transition probabilities with (6) and feed them to (3) to estimate the new expected time. Therefore, once captured the signatures of K grasp executions, the optimization process can be done completely offline and without any extra required experiments.

Since an analytical expression of τ directly as a function of the thresholds $\pi_1 \dots \pi_n$ is not feasible, and thanks to the fact that evaluating τ is fast and simple, we use a direct search gradient-free method to optimize it. In Sect. VII we detail the implementation of the optimization and the results obtained, in particular how the expected time decreases with the number of abort points n.

VI. Implementation

A. System Architecture

The system implementation has a modular design, based on the ROS (Robot Operating System) architecture [19]. Each subsystem is contained within a separate node, with messages being passed between the nodes containing both sensory data and commands.

A finite state machine that implements the Markov chain in Fig. 5 governs the overall system. It cycles through each one of the steps of the grasp and allows for easy modification of the grasp behavior. Different nodes within the system include:

- **Main Controller**: Primary node which implements the state machine.
- **Robot Controller**: Controls the position of the industrial manipulator.
- **Grasp Controller**: Controls the motor in the hand, and broadcasts motor and finger encoder positions.
- **Vision Interface**: Aggregated vision routines to provide ground truth for the learned models both on the number of markers grasped and their position within the hand.
- **Learning Interface**: Receives motor and encoder readings, and broadcasts success probabilities.

B. Vision System

The data-driven approach used to model the probability of success in Sect. VI-C requires running a large number of grasp executions and logging both their signatures and outcomes. The vision system is meant to provide feedback both in terms of the number of objects grasped and the location of the objects within the palm of the gripper, making the overall system self-supervised.

We have implemented a vision system tailored to the specific application and object (highlighter marker) using Willow Garage's OpenCV vision processing library [20]. It is composed of the following steps:

1) **Background subtraction:** We capture an image of the hand with no markers and then black out all areas of the image reasonably similar to the calibration image.
2) **Find color regions:** Since the highlighter markers are brightly colored, we threshold the image to determine regions of color. We then clean those same regions by removing small clusters of color.
3) **Find edges and lines:** We recognize markers by their straight edges using the Canny edge detector near color regions. We then use the Hough line detector to determine prominent lines in the image and assume that the long edges of each marker are among those.
4) **Most likely position for a marker:** Each detected line is scored proportional to the amount of color to each one of its sides. Iteratively we detect the most likely edge of a marker and subtract the color labeled region until insufficient color is left in the image for another marker to exist.

We evaluated the vision system with 266 images captured in successive trials. In the task of classifying the grasp outcome between success and failure, the algorithm was able to correctly classify all images except one where a marker was caught in the unlikely position of pointing perpendicular to the palm. The accuracy of the vision algorithm is high enough to treat its output as ground truth for the posterior learning system.

C. Learning System

As detailed in Sect. III-C, the objective of the learning system is to create a probabilistic model between the sig-

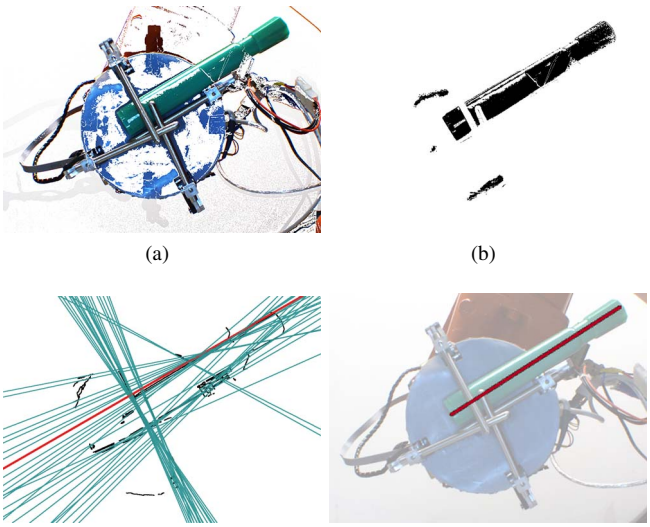

(a) (b) (c) (d)

Fig. 8. The vision system outputs the position and orientation of the marker within the palm of the hand. (a) Filtered image after background subtraction. (b) Image after color region filtering. (c) Edge and line detectors. (d) Most likely position of the marker.

nature of a grasp $\mathcal{G}$ and the success probabilities at some predefined abort points.

Among the different available techniques for probabilistic classification we choose *Relevance Vector Machines* (RVM) [21], which employ a similar formulation as Support Vector Machines, but use Bayesian inference to provide probabilistic classification. We use the implementation provided in the Dlib Machine Learning library [22].

Prior to training the RVMs, we use *Principal Component Analysis* (PCA) [23] to reduce the dimensionality of the grasp signature. At each abort point t_i, we compress the section of the signature $[0, t_i]$. PCA finds a linear transformation of the signature into a smaller number of linearly uncorrelated features while retaining most of the original variability across the set of signatures.

After compression of the signatures, we use half of the training set to learn the RVMs. The other half will be used in Sect. VII to optimize the cut-off probabilities. Figure 4 shows an example of the evolution of the estimation of the success probability provided by the trained RVMs.

VII. Results

In order to optimize the cut-off probabilities $\pi_1 \dots \pi_n$ we first capture the signatures of $K = 200$ grasp executions. Out of those we draw randomly $\frac{K}{2}$ that we use to train the the probabilistic classifiers as detailed in Sect. VI-C. We use then the other $\frac{K}{2}$ to optimize the probability thresholds.

For any given value of the cut-off probabilities, we make use of (3) to efficiently evaluate the expected time to a successful grasp. We then use *ga* optimizer in Matlab, to optimize and the cut-off probabilities.

We normalize all obtained expected times by T, the time span of the grasp, so that $\tau = 1$ is the asymptotically optimal solution. In the case of no early abort, the expected time is $\frac{1}{f}$. The improvement when using early abort is measured as the percentage of decrease of the expected time from that same baseline, $\frac{1}{f}$, towards the optimal.

TABLE I

Normalized expected time to a successful grasp.

n	τ	Improvement
1	2.17	-
2	2.12	4.3%
4	1.98	16.5%
8	1.91	22.0%
16	**1.58**	**50.4%**

Table I details the variation of the normalized expected time to a successful grasp with the number of abort points n after optimization. The case $n = 1$ is the baseline to compare with (system without early abort). Around $n = 16$ the optimization problem gets too big to be addressed by the off-the-shelf optimizer *ga* in Matlab. It is sufficient, though, to demonstrate that early abort reduces the expected time to a successful grasp, in the studied case with an improvement of up to 50%.

VIII. Discussion and Future Work

When looking at humans manipulating the world, one soon realizes that mistakes are not uncommon. As an example, when grasping an object one of the fingers might contact it in the wrong place, the object might slip from the fingertips in an attempt to lift it, or they might tip it over in an attempt to place it. Part of our skill set, however, is being good and fast at recovering from those failures.

Following that same idealistic goal, in this paper we introduce the concept of early abort and retry in the context of grasping in a bin picking task. We allow a hand to abort and retry the grasp as soon as it is confident that it will fail. In doing so, we improve the efficiency of the system with respect to earlier work and allow a simple hand to be competent in solving a complex task.

The main contribution of this paper is to show that we can model a iterative system with a set of predefined abort points as a Markov chain and use the model to optimize the expected time to successful completion of the desired task.

Although we have focused on grasping in a bin-picking scenario, the proposed methodology generalizes to any process generating a signature that correlates with the potential success or failure of the execution. Automated assembly is an example of application that would benefit from early abort to improve their performance.

Our long-term goal is to demonstrate broad manipulation capabilities with simple hands. In earlier work, we approached the bin-picking problem with a blind policy driving the hand. Early abort is a step forward that suggests a binary policy that at each instant allows the hand either to abort or to continue with the execution. In our process to gradually complexify the grasp policy, we intend in future work to learn an optimal singulating policy from a small parametrized set.

REFERENCES

[1] M. T. Mason, S. S. Srinivasa, A. S. Vazquez, and A. Rodriguez, "Generality and Simple Hands," Robotics Institute, Carnegie Mellon University, Technical Report CMU-RI-TR-10-40, 2010.

[2] A. Rodriguez, M. T. Mason, and S. S. Srinivasa, "Manipulation Capabilities with Simple Hands," in *International Symposium on Experimental Robotics (ISER)*, 2010.

[3] R. Tella, J. Birk, and R. Kelley, "General Purpose Hands for Bin-Picking Robots," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 12, no. 6, pp. 828–837, 1982.

[4] K. Boehnke, "Object Localization in Range Data for Robotic Bin-Picking," in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2007, pp. 572–577.

[5] S. Cho, S. Asfour, A. Onar, and N. Kaundinya, "Tool Breakage Detection Using Support Vector Machine Learning in a Milling Process," *International Journal of Machine Tools and Manufacture*, vol. 45, no. 3, pp. 241–249, 2005.

[6] Y.-W. Hsueh and C.-Y. Yang, "Prediction of tool breakage in face milling using support vector machine," *The International Journal of Advanced Manufacturing Technology*, vol. 37, no. 9-10, pp. 872–880, 2007.

[7] D. M. Tax, A. Ypma, and R. P. Duin, "Support Vector Data Description Applied to Machine Vibration Analysis," in *ASCI'99*, vol. 54, no. 1, Heijen, Netherlands, Jan. 1999.

[8] K. Althoefer, B. Lara, Y. H. Zweiri, and L. D. Seneviratne, "Automated Failure Classification for Assembly with Self-Tapping Threaded Fastenings Using Artificial Neural Networks," *Journal of Mechanical Engineering Science*, vol. 222, no. 6, pp. 1081–1095, 2008.

[9] A. Rodriguez, D. Bourne, M. T. Mason, G. F. Rossano, and J. Wang, "Failure Detection in Assembly : Force Signature Analysis," in *IEEE Conference on Automation Science and Engineering (CASE)*, 2010.

[10] A. M. Dollar, L. P. Jentoft, J. H. Gao, and R. D. Howe, "Contact sensing and grasping performance of compliant hands," *Autonomous Robots*, vol. 28, no. 1, pp. 65–75, Aug. 2010.

[11] A. Bicchi, J. K. Salisbury, and D. L. Brock, "Contact Sensing from Force Measurements," *The International Journal of Robotics Research*, vol. 12, no. 3, p. 249, 1993.

[12] D. Siegel, "Finding the Pose of an Object in a Hand," in *IEEE International Conference on Robotics and Automation*, 1991, pp. 406–411.

[13] Y. Jia and M. A. Erdmann, "Geometric Sensing of Known Planar Shapes," *The International Journal of Robotics Research*, vol. 15, no. 3, pp. 365–392, 1996.

[14] ——, "Pose and Motion from Contact," *International Journal of Robotics Research*, vol. 18, no. 5, pp. 466–490, 1999.

[15] A. S. Wallack and J. F. Canny, "Generalized Polyhedral Object Recognition and Localization Using Crossbeam Sensing," *The International Journal of Robotics Research*, vol. 16, no. 4, pp. 473–496, 1997.

[16] J. Laaksonen, V. Kyrki, and D. Kragic, "Evaluation of Feature Representation and Machine Learning Methods in Grasp Stability Learning," in *IEEE International Conference on Humanoid Robots*, 2010, pp. 112–117.

[17] P. Viola and M. J. Jones, "Robust Real-Time Face Detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, May 2004.

[18] M. Dundar and J. Bi, "Joint Optimization of Cascaded Classifiers for Computer Aided Detection," in *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2007, pp. 1–8.

[19] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: an Open-Source Robot Operating System," in *International Conference on Robotics and Automation*. IEEE, 2009.

[20] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[21] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *The Journal of Machine Learning Research*, vol. 1, no. 3, pp. 211–244, Aug. 2001.

[22] D. King, "Dlib-ml: A Machine Learning Toolkit," *The Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.

[23] L. Smith, "A tutorial on principal components analysis," USA, 2002.