Learning to grab objects with the NAO robot

Eric Jansen and Tijn van der Zant

Department of Artificial Intelligence, Faculty of Mathematics and Natural Sciences,
University of Groningen
Nijenborgh 9, 9747 AG, Groningen
{eric.jansen,c.m.van.der.zant}@rug.nl
http://www.rug.nl/ai

Abstract. Using an interval estimation algorithm a NAO robot is learned to select a behavior to pick up an object it is presented with. By updating the success rate of the behaviors that are used, the NAO robot learns online to pick the right behavior. After six trials with various items based on the items in the RoboCup@Home competition, 78% of the items can be picked up without failure.

Keywords: NAO, Robotics, RoboCup@Home, Object manipulation

1 Introduction

With a sharp rise in the ageing population and a general desire people have to avoid tedious work, the demand for a robot that is able to take over household chores is ever rising. A competition that is designed to motivate researchers to build a robot that fits this description is the RoboCup@home competition. A domestic robot has to be able to manipulate objects in order to perform most domestic chores. Borst (Borst, Fischer, & Hirzinger, 1999), among others (Miller, Knoop, Christensen, & Allen, 2003) (Morales, Asfour, Azad, Knoop, & Dillmann, 2006), has shown ways for robots to plan the exact motion their hand has to follow in order to pick up an object. However, planning is time-consuming and doesn't allow for real-time interaction with the world. A database with predefined behaviors also allows for robust grasping of objects (Goldfeder, Ciocarlie, Dang, & Allen, 2009). Unfortunately, such a database is focussed on very complex robot hands. A robot with fewer possibilities of picking up an object cannot benefit from the complex behaviors.

1.1 Research question

By creating a small database of behaviors and letting a small and simple robot learn to select a behavior, this research will focus on grabbing objects in real-time with simple hands and arms. More precisely, the subject of this research is answering the following question:

How well can a NAO robot learn to grab different kind objects, using machine learning and a limited set of behaviors?

'How well' can be defined by the number of attempts and overall success of picking up an object. Since this behavior is meant to be used in the RoboCup@home competition, the items that are used during the experiments are based on items that have been used in previous editions of the competition.

Goal of the experiments
The experiment is done to show that learning can occur fast and the robot is able to cope with varying objects by selecting from several behaviors. The robot is therefore not expected to have a perfect record after the experiments, but it has to be able to vary it's behavior and cope with new objects.

1.2 The robot

The University of Groningen has one participating team in the RoboCup@Home competition, called BORG. The robot they use is a hybrid of a pioneer robot and a NAO robot (see Figure 1), combined with extra sensors, such as a Microsoft Kinect camera, and additional computing power. The arms of the NAO robot are the main actors of the robot with respect to object manipulation. The robot is manufactured by Aldebaran Robotics (*Website of Aldebaran Robotics*, 2012). It is designed to be a robot with educational purposes.

Arms and hands
The short arms and hands resemble those of a human being. However, the robot only has two fingers and a thumb and the fingers cannot be moved separately. Furthermore, the hand is not able to exert a lot of force on something that is grabbed. The arms, however, are stronger and more capable of holding an object tight.

2 Methods

The experiment is focused solely on grabbing objects and selecting the appropriate behavior, so it is assumed that the robot knows what object it is facing. In order to pick up an object, the robot can select his behavior from a list of six different grabbing behaviors. A set of nine objects, differing in size, weight and physical appearance, is selected. The different behaviors and objects are discussed in more detail below.

2.1 Behaviors

The robot always starts from the same starting position (see Figure 1). From there, the robot initiates the grabbing behavior it has selected. Both the starting

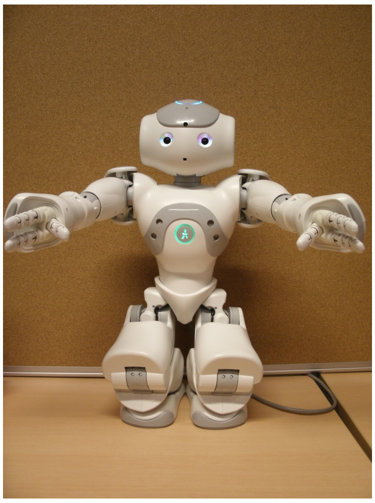

Fig. 1. The NAO robot by Aldebaran Robotics, currently in its starting positions prior to picking up an object

position and every grabbing behavior has been taught to the robot by recording it via a part of the software that comes with the robot: Choregraphe. The recorded behaviors were inspired by natural behaviors primates exibit (Leca, Gunst, & Huffman, 2010).
There are six behaviors the NAO can choose from. Since two are mirrored implementations of others, four descriptions are needed to explain all the behaviors:

1. **Two-handed grab** The robot moves both hands towards the object, clenching it between both hands. Note that the force that keeps the object in place comes from the arms, not from the hands.
2. **One-armed chest grab (1 & 2)** The robot leaves one arm unmoved. The hand of the other arm is moved towards the shoulder of the unmoved arm. An object can now be held between the moved arm and chest of the robot. This describes two different behaviors, because this can be done with either arm.
3. **Two-armed chest grab (1 & 2)** The robot lowers one arm and moves both arms towards his chest, resulting in a hug-like stance. The robot holds on to the object by keeping it pressed to its chest. Since either arm can be lowered, this description also applies to two separate behaviors.
4. **Two-handed lift** The robot moves its arms to sides of its body and points its forearms forward and slightly downward. The hands are moved upward by turning the shoulders upward. If an object has the physical properties that allow for the robot to place its hands underneath the object, the robot can lift the object.

See Figure 2 for the final position of the robot at the end of each behavior.

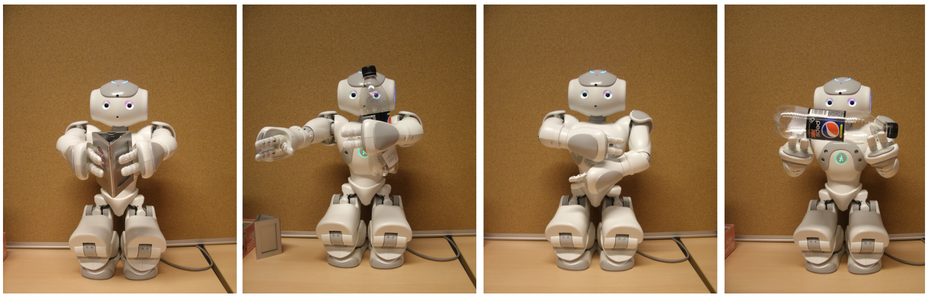

Fig. 2. The four final positions after the behaviors executed. Note that both the second and third position have a mirrored version.

2.2 Objects

The set of objects is based on objects that have been used in previous editions of the RoboCup@Home competition. That means they can all be found in a typical household in or around the living room and kitchen.

Name	**Description**	**Size (cm)**	**Weight (grams)**
Ball	A standard tennis ball	⌀6.5	57
BottleP	An empty pepsi-bottle, 0.5 liters	⌀6, ↑23	28
Box	Lid of a game box	18 × 11 × 7	46
Can	An empty Coca-Cola can, 0.33 liters	⌀6.5, ↑11.5	25
Cup	An empty, plastic coffee cup in a small holder with extra handle	⌀5.75, ↑8.5	8
DessertCup	An empty cup for chocolate mousse	⌀8, ↑5	8
Pringles	A small, empty Pringles can, original content: 40 grams	⌀8, ↑8.5	25
Sponge	A standard scourer, dry	18 × 11 × 7	46

2.3 Learning algorithm

In order to select a behavior for picking up a presented object, an implementation of the interval estimation algorithm is used (Zant, Wiering, & Eijck, 2005). The algorithm decides what behavior to select by estimating the success of grabbing the specific object with each behavior. Because the trials can be seen as Bernoulli trials, the success is binomially distributed. This distribution can be approximated by a normal distribution. The alorithm goes through the following steps:

1. Calculate the distributions of the success rate for every behavior for the object that is presented;
2. Select the behavior that has the confidence interval with the highest upper bound. If two or more behaviors have an equally high upper bound, select the first;
3. Execute the selected behavior;
4. Record whether or not the attempt was successful.

2.4 Experimental setup

For the experiments, action-selection was turned off. The robot used the same behavior six times on the same object before moving on to the next behavior. This way, the robot generated its own estimations regarding objects per behavior. The object was placed on a small table and the robot was positioned in front of it. Visual feedback was not a part of this experiment. Feedback on success was therefore given via input of a supervisor. Success was determined by removing the table on which the object rested. If the robot held on to the object, the trial was considered to be successful.
Objects were roughly placed on the same in every trial, but it was never exactly the same, emulating the real world, where the positioning of the robor

3 Results

3.1 Grab success

After six trials per object for every behavior, seven out of nine objects could be picked up every time. The other two objects had an 83% and a 67% chance, respectively, of being picked up. See Figure 3.

3.2 Learning

Using the upper bound of the confidence intervals to determine what behavior to use, the NAO learned what behavior to use for what object in six trials. This is not enough to statistically prove it always chooses the best behavior. However, it does show that the NAO can learn to pick up a new object quite fast.
The upper bound of the confidence interval is subject to fluctuation, so if a trial fails, the NAO can recover fast and select a different behavior to try again.

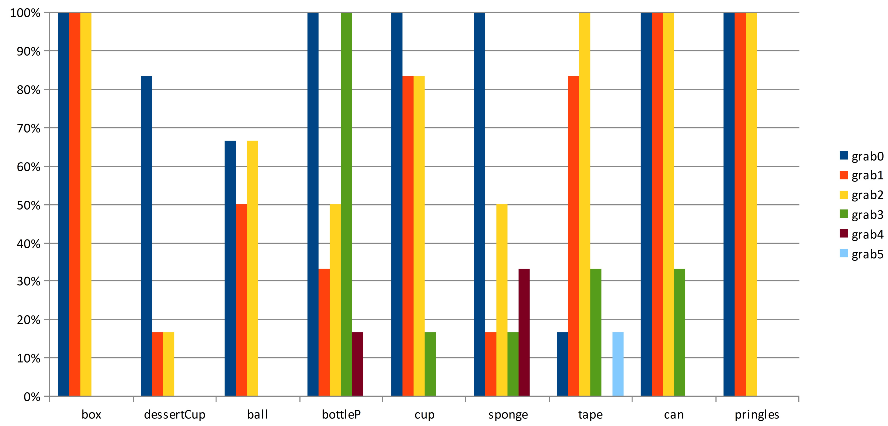

Fig. 3. The success rate of picking up each object with every behavior after six trials

4 Discussion

The goal of this project was to use a NAO robot to pick up several objects that reflect objects used in the RoboCup@Home competition. After only six trials, the NAO was able to pick up seven out of nine objects every time. The remaining two objects had no perfect pick up rate, but the probability of being picked up was still high, at least 67%.
Because of the scope of this research, not enough trials were conducted to have statistical evidence that this method of learning behaviors actually works well. However, the tendency of the results are promising. Behaviors can be switched easily if necessary and most objects had a high success rate of being picked up. The amount of behaviors was held small, but the framework allows for extending this number. If the NAO has to pick up new objects, the same experimental setup can be used to learn what behavior to use.

4.1 Future work

In this paper, only behaviors were used that were manualy created for the NAO. By generalizing the set of behaviors or making the NAO learn the behaviors for itself, the system will be better generalizable.
Furthermore, the set of objects was chosen to mimic objects that have been used in the RoboCup@Home competition. To get a more general pick up robot, the objects that are used, the variety of objects should be higher.

5 References

References

Borst, C., Fischer, M., & Hirzinger, G. (1999). A fast and robust grasp planner for arbitrary 3d objects. *IEEE International Conference on Robotics and Automation, 1999. Proceedings.*, *3*, 1890 - 1896.

Goldfeder, C., Ciocarlie, M., Dang, H., & Allen, P. K. (2009). The columbia grasp database. *IEEE International Conference on Robotics and Automation, 2009. Proceedings.*, 1710 - 1716.

Leca, J.-B., Gunst, N., & Huffman, M. A. (2010). Principles and levels of laterality in unimanual and bimanual stone handling patterns by japanese macaques. *Journal of Human Evolution*, *58*(2), 155 - 165. Available from `http://www.sciencedirect.com/science/article/pii/S0047248409001717`

Miller, A. T., Knoop, S., Christensen, H. I., & Allen, P. K. (2003). Automatic grasp planning using shape primitives. *IEEE International Conference on Robotics and Automation, 2003. Proceedings.*, *2*, 1824 - 1829.

Morales, A., Asfour, T., Azad, P., Knoop, S., & Dillmann, R. (2006). Integrated grasp planning and visual object localization for a humanoid robot with five-fingered hands. *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2006*, 5663 - 5668.

Website of aldebaran robotics. (2012, June). Available from `http://www.aldebaran-robotics.com/en/`

Zant, T. v. d., Wiering, M., & Eijck, J. v. (2005). On-line robot learning using the interval estimation algorithm. *Proceedings of the 7th European Workshop on Reinforcement Learning*, 11 - 12.