

Humanoid Robots: Storm, Rogue, and Beast

Jacky Baltes, Chi Tai Cheng, Stela Seo, and John Anderson

Autonomous Agent Lab
University of Manitoba
Winnipeg, Manitoba
Canada, R3T 2N2

j.baltes@cs.umanitoba.ca
<http://www.cs.umanitoba.ca/~jacky>

Abstract. This paper describes our latest humanoid robots STORM, ROGUE, and BEAST. Similar to previous years, all of our latest generation of humanoid robots are based on the commercially available Bioloid Robotics kit. The Bioloid kit provides a mechanically sound and affordable platform, but does not provide facilities for on-board computer vision and other sensors for active balancing. Thus, it is not suitable as a research platform for humanoid robotics. To overcome these limitations, we added a mount for a Nokia 5500 mobile phone which serves as brain of the robot: it provides computer vision, active balancing, and higher level reasoning. The Nokia 5500 communicates with the AVR AtMega128 micro-controller via a custom designed IRDA infrared. We developed software to create, store, and play back motion sequences on the Robotics Bioloid as well as a new inverse kinematics engine.

1 Introduction

Humanoid robots have always inspired the imagination of robotics researchers as well as the general public. Up until 2000, the design and construction of a humanoid robot was very expensive and limited to a few well funded research labs and companies (e.g., Honda Asimov, Fujitsu HOAP). Starting in about 2001 advances in material sciences, motors, batteries, sensors, and the continuing increase in processing power available to embedded systems developers has led to a new generation of affordable small humanoid robots (some examples include: Pino [7], Manus I [8], Tao-Pie-Pie [2], Roboerectus [9], and Hansa Ram [5]).

The creation of these humanoid robots also coincided with an increased interest in several high profile research oriented international robotics competitions (e.g., RoboCup [3] and FIRA [1]). The researchers chose robotic soccer as a challenge problem for the academic fields of artificial intelligence and robotics. Robotic soccer requires a large amount of intelligence at various levels of abstraction (e.g., offensive vs defensive strategy, role assignment, path planning, localization, computer vision, motion control). Robotic soccer is a dynamic real-time environment with multiple agents and active opponents that try to prevent the robot from achieving its goal. These competitions allowed researchers to compare

their results to others in a real-world environment. It also meant that robustness, flexibility, and adaptability became more important since these robots had to perform for extended periods of time in variable conditions. This is in contrast to researchers that could previously fine tune their system to the specific conditions in their laboratory. The inaugural humanoid robotics competition at RoboCup and at FIRA were held in 2002.

The following section describes the hardware of the Robotics Bioloid robot and our modifications. The vision processing part of our system is described in section 5. Section 6 describes our pragmatic AI method for localizing the robot in the playing field and mapping the environment around the robot. The paper concludes with section 8, which also gives directions for future work.

2 Team Members

For the last three years, the UofM Humanoids team is an important and integral part of our research into artificial intelligence, computer vision and machine learning. Various students and staff have contributed to the 2011 team and a comprehensive list would be too long. The following table lists the core team members of the UofM Humanoids 2011.

Jacky Baltes	team leader, hardware	John Anderson	robot coordination
Chi Tai Cheng	electronics	Stela Seo	motion planning
Yan Wang	vision processing	M.C. Lau	human robot interaction

3 Hardware Description

The Robotis Bioloid robot is a humanoid robotics kit with 19 degrees of freedom (DOF) controlled by serial servo motors. These include two in each ankle for frontal and lateral movement of the foot, one in each knee, three at each hip for frontal, lateral, and transversal movement of the leg, three in each arm, and one to tilt the head. The AX12+ servo provides 10kg/cm at 10V of torque and a maximum speed of 0.17sec/60°. These servos are powerful given their cost and were the main reason that we decided on this robot kit.

Even though the Bioloid kit is a versatile and robust base for humanoid robotics it lacks sensors and processing power to make it suitable for humanoid robotics research. We therefore added a Nokia 5500 mobile phone to the kit as well as a custom built IRDA interface to the basic kit. This setup is the basic for our humanoid robots STORM, ROGUE, and BEAST. Figure 1 shows a picture of our robot STORM.

The AX12+ servos are controlled by a AVR AtMega128 based micro-controller board. Robotis provides a simple firmware for the AtMega128 which is able to record and play back motions. However, this firmware is unsuitable for the challenges of robotic soccer. We therefore implemented our own firmware which required reverse engineering part of the original Robotis firmware (e.g., the control



Fig. 1. The modified Robotis Bioloid robots STORM, ROGUE, and BEAST

of the charging circuit for the battery). We added support for 3 axis accelerometers from Analog Devices to the firmware as well as the ability to merge motions.

The AX12+ servos use a serial protocol, which greatly simplifies the cabling of the robot. In practice the number of servos is limited by the maximum current that can be driven by the controller board.

The ATmega128 embedded system accepts commands via a serial line. The controller boards accept commands at 115200 baud. Commands include setting the position (i.e., set point for all 18 servos), writing a motion (i.e., a sequence of positions), as well as reading back positions and motions. It also includes commands to receive current accelerometer readings from the phone, which are multiplied by correction matrices to actively balance the robot.

The problem is that mobile phones do provide very limited IO resources. We therefore designed our own IRDA interface circuit based on the Microchip MCP 2150 IrDA transceiver.

After these modifications, STORM has a height of 41cm, each foot has a surface contact area of $41.85cm^2$, and the center of gravity is at a height of 18.5cm.

4 Motion Development

Figure 2 shows the interface of the motion control software that we developed to control STORM, ROGUE, and BEAST. The interface allows one to move the robot into a specific position and save this position. The interface also allows one to set the trim (i.e., offset) for all joints as well as the home position.

A separate window tab is used to combine positions into motions. Each motion has a cycle time associated with it and each part of a motion has a arrival time associated with it. Thus, the interface allows the user to easily adjust the speed of a whole motion or individual parts of the motion. The trajectory of all joints is shown in the bottom window.





Fig. 2. Interface of our motion development system. The top window shows the development of a position, the bottom window shows the combination of these positions into motions.

About 20 movements were programmed into the robot including: start walking, take step with right foot, take step with left foot, stop from left walk, stop from right walk, sideways step left, sideways step right, and kick with right foot.

These movements are then available to be played back as required by any of our programs running on the mobile phone.

5 Vision Processing

STORM uses a CMOS camera as the main sensor. The camera is used to approach objects in the field of view of the robot as well as localization and mapping.

To be robust enough to deal with the complex environment of robotic soccer, the vision processing makes little use of colours, but uses a very fast approximate region segmentation algorithm. First, the algorithm scans the image and extracts scan line segments (i.e., segments of similar colour) of approximately the right size. This step is similar to standard region segmentation algorithms.

However, we noticed that implementing a full union-find algorithm was too slow since it took about 2 secs. per image. Since most objects of interest in the environment are relatively small, we use a flood fill pixel merge algorithm, to find the associated region for a scanline. Note that the flood fill algorithm keeps track of which pixels have previously been visited and thus will visit each pixel at most once. The returned region is then checked for size (i.e., number of connected pixels), size of the bounding box, aspect ratio, and compactness. Only in the final step does the algorithm test whether the average colour of the region matches the object colour. If any of these tests fail, the object is rejected. Using only average colours of regions results in robust recognition of the ball and the goals and takes on average approximately 200ms.

An approximation of the relative position of objects is possible by determining the pan and tilt angles, and then calculating the distance to the centre of the image. It is assumed that these objects are on the ground plane. The relative position of an object at the centre of the image will have the closest approximation, so the camera is centered on important objects such as the ball before a decision is made as to what action to take next.

Goals are also detected as objects. Each goal is a distinct colour according to RoboCup rules. If both goal colours are found in one image, the regions of each goal colour are merged with other regions of the same goal colour. The goal colour that is present in the largest merged region is considered to be the goal currently being viewed.

To help the feature based localization method described in the following section, we use a complex camera calibration based on the Tsai camera calibration algorithm [6]. This calibration is only done once for each robot. Given this calibration information, we are able to map points in the image accurately to their real world coordinates. This is essential since it allows us to determine the distance and orientation of the ball to a feature point (ball, goal post, line)

Before localization can occur, features must be extracted from the image. The relevant features for localization on the soccer field are lines, goals, and the centre circle. We use the lines and the goals to achieve localization.

Every 5th column, the system scans from the bottom of the image towards the top. If there is a transition from a green pixel to a white pixel, the pixel p is remembered in a list. The scan continues upward, so there may be more than one transition pixel in a column.

Next, lines are found by running a gradient guided Hough transform [4]. For each point p_i , a set of adjacent points is determined. Triplets are formed from these by including one point to the left of the point p_i , and one point to the right of p_i . There are several triplets that can be formed this way out of the neighborhood of adjacent points. Each triplet votes for an unbounded line in the image. This vote is fuzzified by voting for a small range of slopes through the point p_i .

The peaks in the Hough accumulator space determine the equations of possible lines. For each peak in the accumulator space, we search along the pixels determined by the line equation to find start and end points of the lines. This results in a set of line segments.

The line segments are ordered based on their size. The longest line segment is assumed to represent the edge of the playing field. Given the distance and gradient of the line segment, the position and direction of the robot can be computed.

6 Localization and Mapping

Knowing the position of the ball is important. Its relative position from the robot is easily determined from an image. However, without knowing the world position of the ball, the robot would often kick the ball out of bound or even into its own goal. Actions can not be taken toward kicking the ball until its world position is known.

Absolute localization of the robot will give absolute localization of the ball. Absolute localization can be done as long as a point is viewed with a known world coordinate, and knowing the robot's world bearing from it. One instance of this is when a goal post is seen. Once this is accomplished, dead reckoning can be used with some accuracy for a short time afterward.

7 Agent Architecture

This section will give a brief introduction to the behaviour tree based agent architecture used in Abarenbou. The discussion will focus on the state transitions. More complex features of the agent architecture (e.g., behaviour trees and state references) are outside of the scope of this paper.

Designing an architecture that is flexible, versatile, and intuitive enough for an intelligent mobile robot is a difficult problem. This is especially true in the case of autonomous robots with limited processing capabilities.




Fig. 3. Localizing using a known point, its relative position, and relative orientation of a line. Image from the robot with highlighted line segments and the calculated position are shown in the lower image.

Abarenbou uses a behaviour tree to balance the need for deliberate planning and reactive behaviours. The behaviours themselves are implemented as finite state machines. As the complexity of the task increases, the implementation of the state machine becomes more error prone.

We therefore developed a meta language to describe the behaviors in XML. The specification of the behaviors includes preconditions (enter functions) and post conditions (exit functions) of the behaviours. An slightly simplified example of a simple behaviour that scans for a target with increasing sweeps is shown in Table 1.

The XML schemas include additional markup to refer to states by name (`%%State("Random Walk")`) access variables (`%%v`) and to trigger transitions to other states (`%%Transition`).

Behaviours are organized into behaviour trees. Higher level behaviours can override or enable other lower level behaviours. For example, a “Perception” behaviour may disable the scan for target behaviour and enable the state “Target In Front” if it recognizes the target.

One of the design goals of the meta language was to be highly efficient. Therefore, instead of adding a XML parser and interpreter to the agent, the meta language is parsed and interpreted offline and converted into highly efficient C code. This code is then compiled and executed on the mobile phone. For example, the example above shows that the programmer uses state names (e.g., “Random

```

<State id="Scan For Target" >
<Enter>
v(angle) = 0;
if ( previousState == State("Target Right Forward") )
{
    v(newAngle) = 20; /* Turn 20 degrees first */
    v(angleAdjust) = +10;
}
else
{
    %%v(newAngle) = - 20; /* Turn 20 degrees first */
    %%v(angleAdjust) = -10;
}
}
</Enter>
<Process>
if ( ( %%v(newAngle) >= -190 ) &&
    ( %%v(newAngle) <= 190 ) )
{
    if ( %%v(angle) != %%v(newAngle) )
    {
        turn( (%%v(angleAdjust) * TEN_DEGREE) / 10 );
        %%v(angle) = %%v(angle) + %%v(angleAdjust);
    }
    else
    {
        %%v(newAngle) = - %%v(newAngle) - 40;
        %%v(angleAdjust) = - %%v(angleAdjust);
    }
}
}
else
{
    %%Transition("Random Walk");
}
}
</Process>
</State>

```

Table 1. An XML schema for a behaviour that scans for a target by turning right/left with increasing sweeps

Walk,” and “Scan For Target”). However, the states names are converted to integers in the C code.

An additional advantage of using a formalized state representation language with markup in the code is that it is possible to generate other representations. For example, our state compiler automatically generates a state transition graph. Figure 4 shows the state transition graph for a simple approach task. The robot first approaches a target and then walks away from it.




Fig. 4. Automatically generated state transition graph for a simple approach and avoid task. Solid lines are state transitions.

8 Conclusion

This paper describes the modifications and additions that we made to convert the Robotis Bioloid humanoid fighting robot kit into a platform for humanoid robotic soccer.

The addition of a Nokia 5500 provides visual feedback, active balancing sensors, and processing for the robot. The original embedded controller of the Robotis Bioloid robot are used to control the motion. The mobile phone communicates via a custom built IRDA interface with the robot and is able to start/stop the motions of the robot.

We use a vision-based approach to determine the behaviour of the robot. The robot uses lines on the playing field to localize itself on the playing field and to map objects into the robot's environment.

The development of the complex architecture necessary for an intelligent soccer player is simplified through the use of an XML based meta language for behaviour trees. This meta language makes the pre-conditions, process stack, and state transitions explicit. The XML representation is converted into C code, which can be compiled into efficient code and thus does not introduce computational overhead through its use.

References

1. Jacky Baltes and Thomas Bräunl. *HuroSot Laws of the Game*. University of Manitoba, Winnipeg, Canada, May 2004. <http://www.fira.net/hurosot>.
2. Jacky Baltes and Patrick Lam. Design of walking gaits for tao-pie-pie, a small humanoid robot. *Advanced Robotics*, 18(7):713–716, August 2004.
3. RoboCup Federation. Robocup humanoid league 2002. WWW, November 2001. http://www.robocup.org/regulations/humanoid/rule_humanoid.htm.
4. K. Y. Hough. Method and means for recognizing complex patterns. U.S. Patent 3069654, 1962.
5. Jong-Hwan Kim, Dong-Han Kim, Yong-Jae Kim, Kui-Hong Park, Jae-Ho Park, Choon-Kyoung Moon, Jee-Hwan Ryu, Kiam Tian Seow, and Kyoung-Chul Koh. Humanoid robot hansaram: Recent progress and developments. *JACIII*, 8(1):45–55, 2004.
6. Roger Y. Tsai. An efficient and accurate camera calibration technique for 3d machine vision. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 364–374, Miami Beach, FL, 1986.
7. Fuminori Yamasaki, Tatsuya Matsui, Takahiro Miyashita, , and Hiroaki Kitano. Pino the humanoid: A basic architecture. In Peter Stone, Tucker Balch, and Gerhard Kraetschmar, editors, *RoboCup-2000: Robot Soccer World Cup IV*, pages 269–278. Springer Verlag, Berlin, 2001.
8. Ruixiang Zhang, Prahlad Vadakkepat, Chee-Meng Chew, and Janesh Janardhanan. Mechanical design and control system configuration of a humanoid robot. In *Proc. of 2nd Int. Conf. on Computational Intelligence, Robotics and Autonomous Systems (CIRAS 2003)*, Singapore, 15 - 18 December 2003.
9. Changjiu Zhou and Pik Kong Yue. Robo-erectus: a low-cost autonomous humanoid soccer robot. *Advanced Robotics*, 18(7):717–720, 2004.