

Local Movement Control with Neural Networks on Autonomous Robots

Steffen Prüter and Ralf Salomon

Abstract— In the RoboCup small-size league, the positions of the robots are commonly calculated via a camera mounted above the field and different kinds of artificial intelligence methods hosted on a PC outside the field. In this case, the robot’s position must be predicted [2] because of the various time delays until the position data [7] arrived at the robot. This paper focuses on the use of local sensors attached to the robot, and a neural network to estimate the actual robot position. This paper shows how local sensors can compensate for the effect of latency times and how the actual robot position can be estimated. Slip and friction effects that cannot be measured with local sensors can be adjusted by a neural network, for which the implementation is adapted to the low resources available on the robot.

I. INTRODUCTION

In the RoboCup Soccer [1], two robot teams play soccer against each other in different leagues. The small-size league is of particular interest, because it combines engineering tasks, such as building robot hardware and designing electronic components, with computer science applications, such as localization of objects, finding the robots’ positions, and calculating the best path through obstacles. Another interesting challenge emerges from the requirement that all team members have to communicate with each other in order to develop a consistent cooperative behavior. Research on artificial intelligence may help find the optimal solution in all of these areas.




Fig. 1. The physical setup in the RoboCup small-size league.

Fig. 1 illustrates the general setup used by all teams of the small-size league. Two cameras mounted approximately four meters above the floor observe the field of four by five

Steffen Prüter is with the University of Rostock, Faculty of Computer Science and Electrical Engineering, Institute of Applied Microelectronics and Computer Engineering, 18051 Rostock, Germany (email: steffen.prueter@uni-rostock.de)

Ralf Salomon is with the University of Rostock, Faculty of Computer Science and Electrical Engineering, Institute of Applied Microelectronics and Computer Engineering, 18051 Rostock, Germany (email: ralf.salomon@uni-rostock.de)




Fig. 2. The image processing system consists of five stages which all contribute to processing delays also known as latency time.

meters in size on which two teams consisting of five robots play against each other. The processing sequence starting at the camera image and ending at the robots executing their received (action) commands suffers from significant time delays, as Fig. 2 illustrates.

These time delays have the consequence that when receiving a command, the robot’s current position does not correspond to its position in the camera image. Consequently, the actions are either inaccurate or may even lead to improper behavior in the extreme case. For example, the robot may try to kick the ball even though it is not in reach anymore.

Section II discusses how the position correction can be further improved on the robot itself and how local sensors can alleviate this problem to a large extent.




Fig. 3. An omnidirectional drive with its calculation model while the robot is driving in direction of F_3 and wheel 2 and 3 compensate the rotation with F_{1y} equal to $-F_{2y}$, and $|F_1|+|F_2|=-|F_3|$.

Fig. 3 shows the omnidirectional drive commonly used by most robots of the small-size league. As can be seen, an omnidirectional drive consists of three wheels, which are twisted to each other by 120 degrees. This drive has the advantage that a robot can be simultaneously doing both moving forward and spinning around its own central axis. Furthermore, the particular wheels, as shown on the left-hand-side of Fig. 3, yield high grip in the rotation direction, but almost-vanishing friction perpendicular to it. The specific orientation of all three wheels, as illustrated on the right-hand-side of Fig. 3, requires advanced controllers and

exhibit higher friction than standard two-wheel drives. The later requires sophisticated servo loops and PID (PID is the abbreviation of proportional-integrate-differential, for further details please see [9]) controllers. Depending on the carpet and the resulting wheel-to-carpet friction, one or more wheels may slip. As a consequence, the robot leaves its precalculated moving path. To this end, a robot employs its own back-propagation network to learn its own specific slip and friction effects. Section 0 concludes this paper with a brief discussion including possible future research.

II. LOCAL SENSORS

As has been outlined in the introduction, the latency caused by the imaging-processing-and-action-generation loop leads to non-matching robot positions. As a measurable effect, the robot starts oscillating, turning around the target position, missing the ball, etc. An approach to solve the latency problem is to do the compensation calculation on the robot itself. The main advantage of this approach is that the robot's wheel encoders can be used to obtain additional information about the robot's actual behavior. However, since the wheel encoders measure only the wheel rotations, they cannot sense any slip or friction effects directly.

A. Latency time

RoboCup robots are real-world vehicles rather than simulated objects. Therefore, all algorithms have to account for physical effects, such as inertia and delays, and have to meet real-time constraints [3]. Because of the real-time constraints, perfectly exact algorithms would usually require too much calculation time. Therefore, the designer has to find a good compromise between computational demands and the precision of the results [6]. In other words, fast algorithms with a just decent precision are the method of choice here [8].

The latency is caused by various components including the camera's image grabber, the image compression algorithm, the serial transmission over the wire, the image processing software, and the final transmission of the commands to the robots by means of the DECT modules. Even though the system already uses the compressed YUV411 image format [5], the image processing software as well as the DECT modules are the most significant parts with a total time delay of about 200ms. For the top-level control software, which is responsible for the coordination of all team members, all time delays appear as a constant-time lag element. The consequences of the latency problem are further illustrated in Fig. 4 and Fig. 5.




Fig. 4. Due to the latency problem, the robot receives its commands at time t_2 , which actually correspond to the image at time t_0 .

Fig. 4 illustrates the various processing stages and corresponding robot positions. At time t_0 , the camera takes an image with the robot being on the left-hand-side. At the end of the image analysis (with the robot being at the old position), the robot has already advanced to the middle position, missing the ball, etc. An approach to solve the latency problem is to do the compensation calculation on the robot itself. The main advantage of this approach is that the robot's wheel encoders can be used to obtain additional information about the robot's actual behavior. However, since the wheel encoders measure only the wheel rotations, they cannot sense any slip or friction effects directly.




Fig. 5. Problem of stopping the robot at a desired point.

B. Experimental Analysis

In order to compensate for the effects discussed above, the knowledge of the exact latency time is very important. The overall latency time was determined by the following experiment:

The test software was continuously sending a sinusoidal drive signal to the robot [10]. With this approach, the robot drives 40cm forward and then 40cm backwards. Then, the actual robot position as was seen in the image data was correlated with the control commands. As Fig. 6 shows, the duration of the latency time is seven time slots in length, which totals up to 234ms with 30 frames send by the camera.

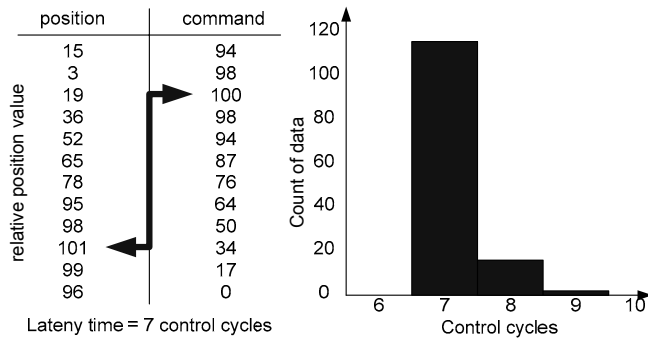


Fig. 6. Detection of the Latency time in the control loop.

It might be worthwhile to mention here that for technical reasons, the time delay of the DECT modules is not constant; the jitter is in the order of up to 8ms. The values given above are averages over 100 measurements.

C. Improved Position Accuracy by Local Sensors

In the ideal case of slip-free motion, the robot can extrapolate its current position by combining the position delivered by the image processing system, the duration of the entire time delay, and the traveled distance as reported by the wheel encoders. In other words: in case slip does not occur, the robot can compensate for all the delays by storing the previous and current wheel tick counts. This calculation is illustrated in Fig. 7.




Fig. 7. Extrapolation of the robot's position using the image processing system and the robot's previous tick count.

Since the soccer robots are real-world entities, they also have to account for slip and friction, which are among other things, nonlinear and stochastic by nature. The following section employs back-propagation networks to account for those effects.

III. EMBEDDED BACK-PROPAGATION NETWORK

Due to the resource limitations of the robot hardware, the number of both nodes and connections that the robot can store on its hardware is quite limited. From a hardware point of view, the memory available on the robot itself is *the* major bottleneck. In addition to the actual learning problem, this section also addresses the challenge of finding a good compromise between the network's complexity and its processing accuracy.

A second constraint to be taken into account concerns the update mechanism of the back-propagation learning algorithm. As is well known, back-propagation temporarily stores the calculated error sums as well as all the weight changes Δw_{ij} [4]. This leads to a doubling of the memory

requirements, which would exhaust the robot's onboard memory size even for moderately sized networks. As a workaround, the approach presented here stores those values on the central control PC and communicates the weight changes by means of the wireless communication facility. This separation is illustrated in Fig. 8.




Fig. 8. Separation of the actual feed-forward network (indicated by FFN in the figure) and the back-propagation training algorithm.

A. Methods

As has been discussed above, the neural network has to estimate the robot position even if slip and/or friction occur. Since the coding of the present problem is but trivial, this section provides a detailed description of it. In order to avoid a combinatorial explosion, the robot is set at the origin of the coordinate system in every iteration. All other values, such as the target position and orientation, are relative to that point.

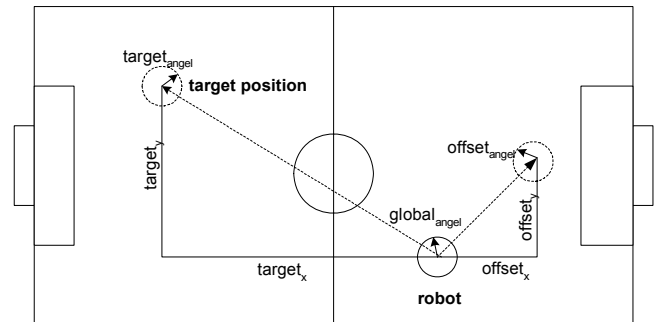


Fig. 9. An example configuration for the slip and friction compensation. For details please see text.

Fig. 9 illustrates an example configuration. This configuration considers three robot positions labeled "global", "offset", and "target". The first robot corresponds to the position as provided by the image processing system. The second position, called "offset", corresponds to the robot's true position and hence includes the traveled distance during the time delay. The third robot symbolizes the robot's target position. As has already been mentioned above, the neural network estimates the robot's true positions (labeled by "offset") from the target position, the robot's previous position, and its traveled distances. The relative values mentioned above are scaled such that they fit into -40 to 40, and all angles are directly coded between 0 and 359 degrees. With all these values the input layer has to have seven nodes.

Due to the limited calculation capabilities of the microcontroller, all values of the neural network may be stored in integer quantities. In this format every operation on the microcontroller is done in two processing steps because of the mathematical coprocessor. For this the feed forward

network calculation (FFN) on the robot must be adapted. To this end a simulation of different FFNs on a PC provide important criteria for the implementation on the robot.

Fig. 10 compares the performance of nine different feed-forward neural network architectures. The initial learning rate was set to $\eta = 0.1$. During the course of learning, the learning rate was increased by 2% in case of decreasing error values and decreased by 50% otherwise. It should be noted that in 10% of all experiments, back-propagation got stuck in local optima. These runs were discarded and are not further considered in this paper. Learning was terminated, if no improvement could be achieved over 100 consecutive iterations. All experiments were done with 400 pre-selected training patterns and 800 test patterns.

As can be seen, one and two hidden layer networks provide a similar accuracy. Networks with more hidden layers are not considered here, since the available computational resources are too limited. The varieties between the average errors of different node counts also low.

The results suggest that a network with five hidden neurons in one hidden layer yields the best performance. This network is a good compromise between network accuracy and calculation time.




Fig. 10. Average error in mm of two and one hidden layer FFNs. On the left hand side two hidden layer networks shown with node count first layer : node count second layer. The three measurements on the right hand side display one hidden layer networks.

The next step is the implementation of the selected FFN on the robot. The measurements show that all resulting network weights Δw_{ij} have a range of -10 to 10. The integer variable on the microcontroller has a range between -32,768 and 32,767. So all weights multiplied by 1,000 to fit into the integer range. The input values, i.e., global, offset, and target, position are multiplied by 100. The network input of all nodes calculated with the formula (1) is a weighted sum of all nodes j to which it is connected by weight Δw_{ij} .

$$net_i = \sum_j w_{ij} o_j \quad (1)$$

Hence on the microcontroller, the node's input net_i is by factor $1,000 \cdot 100 = 100,000$ higher than on the PC side. Therefore, the node's input is stored as long integer as a

consequence of the possibly high range. The neural network learning process on the PC side employs floating-point arithmetic to evaluate the sigmoid function and the node output (2).

$$o_i = 1/(1 + e^{-net_i}) \quad (2)$$

This activation function also be realized on in the microcontroller. The calculation of the results is difficult to implement because of the limited calculation time on the robot. A solution to this problem is to store the needed values in a constant predefined array, because no RAM is required for the activation function. The input values of the array have to be multiplied by 10 and the output values by 10,000. To match the net input net_i , and the sum is divided by 10,000. The simulations on the PC show that the net input net_i has a maximum range of -10 to 10 thus the array consists of 201 values. With these modifications the calculation of the FFN is feasible on the microcontroller.

B. Results

The adjustments of the FFN allow for the implementation on a small-size league robot. Fig. 11 compares the average position error in mm without an FFN to that of FFNs simulated on the PC and on the robot. As be seen, can the FFN provides a gain of 50% in accuracy. The error caused by the modification of the FFN on the robot is not larger than 8%.




Fig. 11. Average error in mm without FFN, the simulated average error on the PC, and the average error on the robot.

A second measurement tests the quality of the correction with a driven eight on the field. In this real world test, the robot drives a specific path that looks like a horizontal eight and rotates around his own axis [10]. This test environment shows how precise and fast the robot can drive. The driven figure has a size of three by one meter and is cut into 64 areas. The PC outside the field monitors the robot position during the measurement and set the new area as target position when the robot has reached the area before it.

Table 1 exemplify that the robot's speed has significant increased on the field via the employment of FFN.

TABLE I.
AVERAGE TIME NEEDED TO DRIVE THE TEST FIGURE.

	Robot	Robot with History	Robot with FFN
Time	8,6s	6,1s	5,4s

This results show that local wheel sensors implemented on the robot advance the accuracy of movement control. But wheel sensors cannot measure slip and friction effects. Back-propagation networks can reduce the positioning errors caused by these effects, but most microcontrollers and embedded devices cannot provide the needed calculation power and memory. The adaptation of a FFN to accomplish the hardware limitations on autonomous robots was also successful. The measurements show that only marginal variations, between the common FFN and the adapted version, occur. So this advancement can be used for further implementations and other developments.

All simulations with different FFN structures and the given input parameters have an error of at most 30 mm. On this point further research must start and investigate other learning and self adaptive principals, such as Hebbian learning [4] and implementation of short-cuts [11]. Furthermore, the input data can be expanded to all values of the history, so that the back-propagation network can also include different acceleration data in the calculation.

IV. CONCLUSIONS

The focus of this paper was on the small-size league, in which two teams of five robots each play soccer against each other. Since no human control is allowed, the system has to control the robots in an autonomous way. To this end, a control software analyzes images sent by two cameras and derives appropriate control commands for all team members.

Unfortunately, the image processing system exhibits various time delays at different stages, which leads to erroneous robot behavior. Section II has shown how local sensors compensate for those effects.

The omnidirectional drives used by most research teams exhibit certain inaccuracies due to two physical effects

called slip and friction. Section III has indicated that neural networks are able to significantly improve the robot's behavior with respect to accuracy, drift, and response.

Furthermore, the architectures presented here still require hand-crafted adjustments to some extent. In addition, the resources available on the mobile robots significantly limit the complexity of the employed networks.

First of all, future research will be addressing the problems discussed above. For this goal, the incorporation of short-cuts into the back-propagation networks seems a promising option. Another important aspect will be the development of complex controllers that would fit into the low computational resources provided by the robot's onboard hardware.

REFERENCES

- [1] <http://www.roboocup.org> (Official Website).
- [2] A. Gloye, M. Simon, A. Egorova, F. Wiesel, O. Tenchio, M. Schreiber, S. Behnke, and R. Rojas: Predicting away robot control latency, Technical Report B-08-03, FU-Berlin, June 2003.
- [3] J. Zagal, and J.Ruiz-de-Solar, Learning to Kick the Ball Using Back to Reality, RoboCup-2004: Robot Soccer World Cup VIII, Springer, 2004
- [4] R. Rojas: Neural Networks - A Systematic Introduction, Springer-Verlag, Berlin, 1996.
- [5] J. C. Russ, The Image Processing Handbook, Fourth Edition, CRC Press, 2002, ISBN: 084931142X
- [6] J.C.Alexander and J.H. Maddocks, "On the kinematics of wheeled mobile robots" Autonomous Robot Vehicles, Springer Verlag, pp.5-24, 1990.
- [7] A. Miene, U. Visser, and O. Herzog, Recognition and Prediction of Motion Situations Based on a Qualitative Motion Description, RoboCup-2003: Robot Soccer World Cup VII, Springer, 2003
- [8] R. Balakrishna, and A. Ghosal, "Two dimensional wheeled vehicle kinematics" IEEE Transaction on Robotics and Automation, vol.11, no.1, pp.126-130, 1995
- [9] K. J. Astrom, T. Hagglund, PID Controllers: Theory, Design, and Tuning, International Society for Measurement and Con; 2nd edition, 1995
- [10] A. Gloye, C. Göktekin, A. Egorova, O. Tenchio, and R. Rojas, Learning to Drive and Simulate Autonomous Mobile Robots, RoboCup-2004: Robot Soccer World Cup VIII, Springer, 2004
- [11] T. Bäck, B. Fogel, and Z. Michalewicz, Evolutionarily Computation 1: Basic Algorithms and Operators, Bristol UK: Institute of Physics Publishing, 2000