

# Local Navigation and Docking of an Autonomous Robot Mower Using Reinforcement Learning and Computer Vision

Ali Taghibakhshi

*Department of Mechanical Science  
and Engineering  
The University of Illinois  
at Urbana-Champaign  
Urbana, Illinois, USA  
alit2@illinois.edu*

Nathan Ogden

*John Deere  
Technology Innovation Center  
Champaign, Illinois, USA  
ogdennatahn@johndeere.com*

Matthew West

*Department of Mechanical Science  
and Engineering  
The University of Illinois  
at Urbana-Champaign  
Urbana, Illinois, USA  
mwest@illinois.edu*

**Abstract**—We demonstrate a successful navigation and docking control system for the John Deere Tango autonomous mower, using only a single camera as the input. This vision-only system is of interest because it is inexpensive, simple for production, and requires no external sensing. This is in contrast to existing systems that rely on integrated position sensors and global positioning system (GPS) technologies. To produce our system we combined a state-of-the-art object detection architecture, You Look Only Once (YOLO), with a reinforcement learning (RL) architecture, Double Deep Q-Networks (Double DQN). The object detection network identifies features on the mower and passes its output to the RL network, providing it with a low-dimensional representation that enables rapid and robust training. Finally, the RL network learns how to navigate the machine to the desired spot in a custom simulation environment. When tested on mower hardware, the system is able to dock with centimeter-level accuracy from arbitrary initial locations and orientations.

**Index Terms**—Reinforcement Learning, Deep Q-Learning, Object Detection, YOLO, Mower

## I. INTRODUCTION

The ever-growing field of autonomous vehicles is a research hotbed for applying Artificial Intelligence (AI) and Machine Learning (ML). In the past few years, there have been a wide variety of improvement in autonomous vehicles. These improvements include, but are not limited to, single agent scale tasks such as path planning, lane changing, and self driving, and multi-agent scale tasks such as collision avoidance and lane management. For instance, in [1], authors have investigated an optimal lane change decision model for autonomous vehicles in high-capacity urban roads. In recent studies [2], [3], researchers have introduced a novel lane management in heterogeneous framework with autonomous vehicles.

Recently, many companies have invested efforts in developing ML-aided driving for increased comfort and safety of their vehicles. However, many autonomous driving and navigation systems still rely on sensing modalities such as Laser Range Finder (LRF), Light Detection and Ranging (LIDAR), and GPS, to name but a handful. These systems can be expensive

and may induce further complications in the design of an autonomous vehicle. It is thus desirable to produce systems which use only vision as the control input.

Over the past decade, reinforcement learning techniques and algorithms have been able to solve complicated decision making problems, both in single-agent [4], [5] and multi-agent frameworks [6]–[8]. With the rise of deep reinforcement learning, there have been multiple studies on implementing value-based learning methods, mostly deep Q-learning [5], in the field of autonomous driving. Using a visual encoding, authors in [9] utilized a low dimensional representation of the driving environment as the observation input for model-free RL agents, whose aim is to achieve urban autonomous driving. [10] has proposed a deep RL framework for autonomous driving, where the RL agents observe raw sensor outputs and take driving actions.

This paper focus on the docking procedure for the John Deere Tango robot mower, which is not reliable in it's current form. This mower was designed to dock using a guide loop wire buried underground, which leads the mower toward the charging station by inducing an electric current in the wire, which is then sensed by a detection system in the mower. Docking failures can occur when the wire is misplaced under the ground or there is a bump in the mower's path. This motivated us to investigate a system that it is more robust to variable initial position and environmental conditions.

In this paper, we have used a combination of a supervised and reinforcement learning algorithms to locally navigate the robot mower and orient it toward the docking station. Firstly, using transfer learning, we train a YOLO network to detect two pre-existing markers on the robot mower to provide positioning and orientation information. Secondly, using the output of the object detection network, we train a DQN agent to learn how to move the robot mower to the desired position employing a curriculum training technique. The real-world scenario is shown in Fig. 1.




Fig. 1: Real-world docking scenario. The mower is shown approaching the charging station under the direction of the RL network, which receives inputs from the YOLO network which localizes the markers in the camera feed. Marker 1 is the stop button and marker 2 is the John Deere logo.

## II. BACKGROUND

### A. YOLO

Real time object detection has been a significant achievement of deep convolutions neural networks (CNN) in the field of computer vision. Using residual connections, deep CNNs can extract complex features from the observed image and are highly accurate in localizing different objects [11]. Generally, these networks are trained on multiple object training sets. Given an image, they draw a bounding box, labeled with the object's tag, around it. The R-CNN algorithm [12] is arguably the pioneer object localization algorithm and, ever since its introduction, many other algorithms have been proposed for the purpose of object detection, including modified versions of R-CNN such as fast R-CNN and faster R-CNN [13], [14]. One of the fastest and most accurate object detection algorithms is You Only Look Once (YOLO), which achieves object detection using a fixed-grid regression [15].

### B. Deep Q-Learning

Value learning is a way of approaching reinforcement learning problems. Deep Q-Learning (DQN) is a value-learning algorithm that has demonstrated success on a wide range of tasks, including achieving human-level control in Atari games [5]. The algorithm combines the traditional Q-learning update with neural network function approximation. Similarly to many RL algorithms, the problem is modeled as a discrete-time Markov Decision Process (MDP). At any time, the agent is in a certain state of the environment's state space,  $S = \{s_1, s_1, \dots, s_n\}$  and has some corresponding actions available from environment's action space,  $A = \{a_1, a_1, \dots, a_m\}$ , which influence the transition to the next state. Transitioning from one state to another provides the agent with a reward  $r$ , and the goal of agent is to maximize the sum of its discounted rewards,

$\sum_{t=0}^{\infty} \gamma^t r_t$ , where  $0 < \gamma \leq 1$  is the discount factor. The state-action value function is denoted by  $Q(s, a)$ , and maps a pair of a state and an action to a real number;  $Q : S \times A \rightarrow \mathbb{R}$ . The Q-learning update [16] is

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left( r_t - Q(s_t, a_t) + \gamma \max_{a'} Q(s_{t+1}, a') \right)$$

where  $\alpha$  is the learning rate. Mnih *et al.* [5] used the Q-learning algorithm in conjunction with deep neural networks and a replay buffer to produce the DQN (Deep Q-Network) algorithm. This uses two networks, the primary and target. The primary network is the one that is being updated using Stochastic Gradient Descent (SGD) at every iteration. The target network is the latest copy of the primary network, and it is used for evaluating the action values. The target network is updated once in every  $N \in \mathbb{N}$  iterations to evaluate the action values using a recent version of the primary network. However, the max operator in DQN algorithm is prone to overestimating the state-action value function since it selects the maximum value of the same Q network. To mitigate this function approximation overestimation, one needs to decouple the selection and evaluation tasks. This leads to the Double DQN algorithm [17], which is used in this research:

---

#### Algorithm 1 Double DQN

---

initialize the primary network  $Q_\theta$ , the target network  $Q_{\theta'}$ , the replay buffer  $D$ , and  $\tau \ll 1$

**for** each iteration **do**

**for** each environment step **do**

    observe state  $s_t$  and select  $a_t \sim \pi(s_t, a_t)$

    execute  $a_t$  and observe the next state  $s_{t+1}$

    and reward  $r_t = R(s_t, a_t)$

    store  $(s_t, a_t, r_t, s_{t+1})$  in the replay buffer  $D$

**end for**

**for** each update step **do**

    sample  $e_t = (s_t, a_t, r_t, s_{t+1}) \sim D$

    compute target Q value:

$$Q^*(s_t, a_t) \approx r_t + \gamma Q_{\theta'}(s_{t+1}, \operatorname{argmax}_{a'} Q_\theta(s_{t+1}, a'))$$

    perform gradient descent step on

$$(Q^*(s_t, a_t) - Q_\theta(s_t, a_t))^2$$

    update the target network parameters:  $\theta' \leftarrow \tau\theta + (1-\tau)\theta'$

**end for**

**end for**

---

## III. SIMULATION AND TRAINING

### A. Motivation for Object Detection with RL

The charging station of the mower is viewed by a Logitech c270 webcam, and the aim is to navigate the mower toward the docking station and either stop it at a desired position or to help it dock, using only the vision input. Therefore, the environment of the problem we are trying to solve, the video feed of the camera, is not only very high dimensional, but it is also dependent on where the setup is located, and it varies from yard to yard. Hence, we lower the dimensionality of the

environment and extract key features of the video feed to both improve system robustness and accelerate RL training.

We use the YOLO algorithm to locate bounding boxes around two markers on the mower, one in the front and the other at the back of the top surface of the mower. The output of the YOLO network is then passed to the RL network as input. Accordingly, the RL agent’s observation space of the world is low dimensional and it is easier to train in this space. The fact that the two markers are located at different ends of the mower allows the RL agent to sense the angle that the mower makes with the straight line toward the docking station. This is due to the fact that both bounding boxes will be in the center of the image if the the mower is exactly oriented toward the docking station. Therefore, in the setup designed in this paper, the centers of the bounding boxes in the picture are the information that is passed to the RL network, which outputs linear velocity and steering rate as actions.

### B. RL Simulation Environment

We have designed a simple simulation environment for training the RL agent. The kinematics of the agent in the simulation environment are ODEs driven by the linear velocity and angular steering rate. The simulation environment simulates the motion and computes the view of the markers from the point of view of the camera. The markers are drawn as quadrilaterals in the simulation, one in red and the other in black, and the rest of the simulated image is just a white background, representing the remainder of the environment. The mower initially starts at a uniformly distributed random position,  $(X, Y) \sim U(-0.2, 0.2) \times U(-0.2, 0.2)$  in meters, and orientation,  $\theta \sim U(-30, 30)$  in degrees. Throughout the paper,  $X$  and  $Y$  coordinates correspond to the position of the rear axle of the mower. The environment is shown in Fig. 2. We stop training the agent when the average reward reaches a certain amount, which was set experimentally. Moreover, in training, we manually reset the mower when its  $Y$  component exceeds 1 m. The mower stops when the  $Y$  component of its position exceeds 1 m and the goal is to minimize the  $X$  and  $\theta$  offsets from zero when it stops. This target goal was chosen because it enables the two metallic rods on the front of the physical mower to connect to the metal pads in the charging station.

The DQN agent takes the center positions of the bounding boxes at the last three time-steps as its input. The reason for using the data at the last three time-steps is to provide the agent with some information about the past to allow the estimation of velocity and acceleration. The agent outputs the desired linear velocity and angular rate based on its observation. The DQN network is shown in Fig. 3, where the state component has two fully connected layers with 16 and 32 neurons and the action component has a single layer with 32 neurons.

For training the agent, we used curriculum training with four phases, numbered  $i = 1, \dots, 4$ . In the first phase the agent is rewarded to go forward, with a small reward for arriving near the target docking location. The subsequent phases increasingly reward more accurate docking. The four

phases are distinguished by different reward functions  $r_t^i$  defined by

$$r_t^i = \begin{cases} R^i & \text{if } Y \geq 1 \text{ m and } |u^1|, |u^2| < c_1^i, \\ 0 & \text{if } Y \geq 1 \text{ m and } |u^1|, |u^2| < 2c_1^i, \\ -(10|v^1(t) - v_0^1| + 10|v^2(t) - v_0^2| + c_2^i|u^1| + c_2^i|u^2|) & \text{otherwise,} \end{cases}$$

where  $(c_1^i, c_2^i)$  are constants with values  $(0.05, 2)$ ,  $(0.05, 5)$ ,  $(0.02, 5)$ ,  $(0.02, 10)$  for  $i = \{1, 2, 3, 4\}$ , respectively;  $R^i = 0$  for  $i = 1$ , and  $R^i = 150$  for  $i = \{2, 3, 4\}$ ;  $u^1, v^1, u^2, v^2 \in (-1, 1)$  are the normalized coordinates of the center of the two markers;  $v_0^1, v_0^2$  are the  $y$  components of the markers when the mower is at  $Y = 1$  m,  $X = 0$  m,  $\theta = 0^\circ$ . The training data is shown in Fig. 4.

### C. Object Detection Network

There are two YOLO object detection networks used in this study, one for detecting the markers in the simulation and another for detecting the real-world markers from the actual camera images. For each of the networks, a pre-trained YOLO network was further trained to detect the markers. The training data for the networks is shown in Fig. 5. Each of the networks was trained on about 3000 labelled images of the markers. The pre-trained feature extraction used ResNet50 [11], and the architecture is re-trained after the 20th ReLU function.

## IV. HARDWARE EXPERIMENTS

### A. Experiment Setup

A Mosquitto MQTT broker was used to send the agent’s actions to the mower via an Ethernet cable, and the connection between the computer and the mower was controlled by a Raspberry Pi. As in the simulation, the mower started from an initial position  $(X, Y)$  and orientation  $\theta$ . The actions were




Fig. 2: Simulation environment. Left: the view from above of the mower approaching the charging station. Right: the simulated camera view of the two markers.




Fig. 3: The RL network architecture. The observation head receives labeled markers in the last three time-steps of the environment and feeds it through the network. The action head also passes each action through a single layer. The last layers of each head are added together and passed through the final layer to output the  $Q$  values.

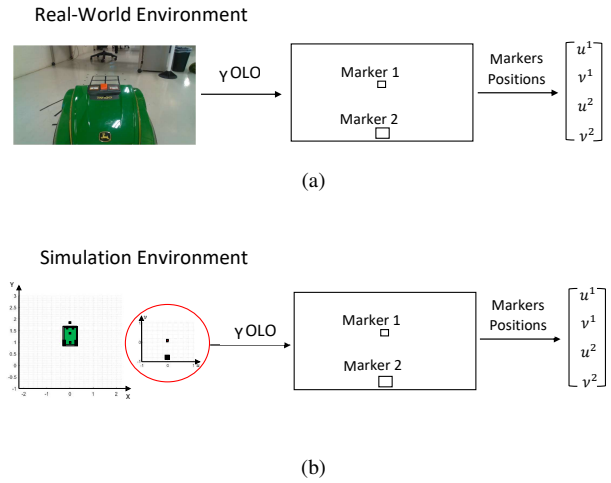


Fig. 5: The two YOLO networks. a) YOLO network trained on real-world mower data. b) YOLO network trained on simulation images. In both cases the image is sent to the corresponding YOLO network and the outputs are the bounding boxes. Finally, the normalized positions of the center of the bounding boxes are extracted as a 4 dimensional vector.

TABLE I: Experiment results.

Error Measure	$X$ Offset (cm)	$Y$ Offset (cm)	$\theta$ Offset (deg)
Max Abs. Error	3.800	3.642	6.200
Mean Abs. Error	0.822	0.934	1.533
RMSE	0.896	1.182	1.661

based the  $X$ ,  $Y$ , and  $\theta$  offsets of the mower when it stopped (when its observed  $Y$  component has exceeded 1 m). The experiment data is shown in Fig. 6 and summary results are given in Tab. I, giving maximum absolute error, mean absolute error (MAE), and root mean squared error (RMSE). The mower had a maximum final position error of less than 4 cm in both  $X$  and  $Y$  directions and a maximum final orientation error of less than  $7^\circ$ , representing successful positioning of the mower in all cases. The mean absolute error was less than 1 cm in both  $X$  and  $Y$  directions and less than  $2^\circ$  in orientation.

## V. CONCLUSION

We demonstrated a cheap and effective control system for autonomous docking of a robotic lawn mower (the John Deere Tango mower), using only vision from a single camera as the sensor. This system was shown to be robust in hardware tests, achieving centimeter-level docking precision.

The controller was a neural network trained using reinforcement learning (Double DQN) in a simple simulated environment. To avoid the need to simulate realistic vision inputs, we trained an object detection network (YOLO) to isolate two markers on the mower. The location of these

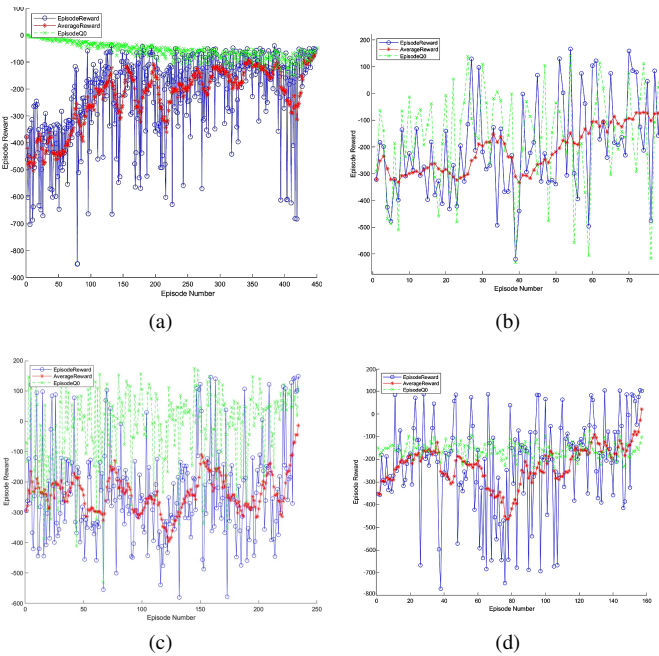


Fig. 4: Training graphs of the agent during the four curriculum phases, with panels a–d corresponding to reward functions  $r^1$  through  $r^4$ .

sent to the mower at a frequency of 5 Hz. The only input to the controlling RL agent was from the camera.

### B. Experiment Results

A total of 90 tests were performed with the initial position and orientation of the mower ( $X, Y, \theta$ ) taken as all combinations of  $X = -0.2, 0, 0.2$  m,  $Y = -0.2, 0, 0.2$  m, and  $\theta = -30, -15, 0, 15, 30^\circ$ , with each initial condition being tested twice. The performance of the agent was measured

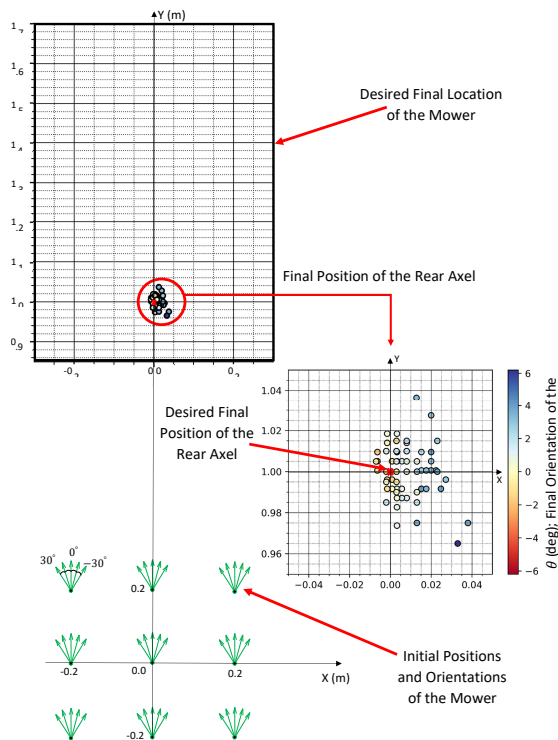


Fig. 6: Experiment configurations. The mower started from one of the positions and orientations shown in green, and from each initial configuration two tests were performed. The desired final position is marked by a red cross. The final positions of the mower are magnified and the final  $X$ ,  $Y$ , and  $\theta$  values are shown.

markers was the only input to the controller agent, making it easy to produce similar inputs in simulation.

In addition to the ease of training, the use of an initial object detection network made the final system robust to different backgrounds and other environment variations. Our choice of markers was opportunistic (we used existing features on the mower) and it is likely that custom markers may be even better. We believe that this paradigm of an object detection network coupled with an RL agent could be an effective strategy for other robot motion control tasks.

#### ACKNOWLEDGMENT

This research was supported by the John Deere Technology Innovation Center.

#### REFERENCES

- [1] P. Cao, Y. Hu, T. Miwa, Y. Wakita, T. Morikawa, and X. Liu, "an optimal mandatory lane change decision model for autonomous vehicles in urban arterials," *Journal of Intelligent Transportation Systems*, vol. 21, no. 4, pp. 271–284, 2017.
- [2] P. Karimi Shahri, S. Chintamani Shindgikar, B. HomChaudhuri, and A. H. Ghasemi, "optimal lane management in heterogeneous traffic network," in *Dynamic Systems and Control Conference*, vol. 59162. American Society of Mechanical Engineers, 2019, p. V003T18A003.
- [3] P. K. Shahri, A. H. Ghasemi, and V. Izadi, "optimal lane management in heterogeneous traffic network using extremum seeking approach," SAE Technical Paper, Tech. Rep., 2020.

- [4] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [5] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [6] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev *et al.*, "grand-master level in starcraft ii using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [7] A. Shojaeighadikolaei, A. Ghasemi, K. R. Jones, A. G. Bardas, M. Hashemi, and R. Ahmadi, "demand responsive dynamic pricing framework for prosumer dominated microgrids using multiagent reinforcement learning," *arXiv preprint arXiv:2009.10890*, 2020.
- [8] A. Ghasemi, A. Shojaeighadikolaei, K. Jones, M. Hashemi, A. G. Bardas, and R. Ahmadi, "a multi-agent deep reinforcement learning approach for a distributed energy marketplace in smart grids," pp. 1–6, 2020.
- [9] J. Chen, B. Yuan, and M. Tomizuka, "model-free deep reinforcement learning for urban autonomous driving," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 2765–2771.
- [10] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani, "deep reinforcement learning framework for autonomous driving," *Electronic Imaging*, vol. 2017, no. 19, pp. 70–76, 2017.
- [11] K. He, X. Zhang, S. Ren, and J. Sun, "deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [12] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 580–587.
- [13] R. Girshick, "fast r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1440–1448.
- [14] S. Ren, K. He, R. Girshick, and J. Sun, "faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [15] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "you only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [16] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [17] H. Van Hasselt, A. Guez, and D. Silver, "deep reinforcement learning with double q-learning," *arXiv preprint arXiv:1509.06461*, 2015.