This article has been accepted for publication in the 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops.

Embedded Robust Visual Obstacle Detection on Autonomous Lawn Mowers

Mathias Franzius
mathias.franzius@honda-ri.de

Mark Dunn
mark.dunn@honda-ri.de

Nils Einecke
nils.einecke@honda-ri.de

Honda Research Institute Europe, Germany

Roman Dirnberger
roman.dirnberger@de.hrdeu.com

Honda R&D Europe, Germany

Abstract

Currently, the only mass-market service robots are floor cleaners and lawn mowers. Although available for more than 20 years, they mostly lack intelligent functions from modern robot research. In particular, the obstacle detection and avoidance is typically a simple collision detection. In this work, we discuss a prototype autonomous lawn mower with camera-based non-contact obstacle avoidance. We devised a low-cost compact module consisting of color cameras and an ARM-based processing board, which can be added to an autonomous lawn mower with minimal effort. On the software side we implemented a color-based obstacle avoidance and a camera control that can deal with the challenging outdoor lighting conditions. For testing our system, we conducted a field test with 20 prototype units distributed in eight European countries with a total mowing time of 3,494 hours. The results show that our proposed system is able to work without expert interaction for a full season and strongly reduces collision events while still keeping the good mowing performance. Furthermore, a questionnaire with the testers revealed that most people would favor the camera-based mower over a non-camera-based mower.

1. Introduction

Autonomous lawn mowers are on the verge of a major market in the lawn and garden segment. Currently, the segment is still small with an installation volume of 103k units in 2015 [11]. However, autonomous mowers are becoming increasingly popular (market growth +25% in 2015 [11]) due to an increase of leisure time and good cutting results.

Currently, autonomous cleaning robots [6, 10, 23] and autonomous lawn mowers [12, 6] are the most promising entry points for robots at home. While high-end research robots for household tasks are typically too expensive and not robust enough for 24/7 application, autonomous cleaning and autonomous lawn mowers have become a robust and stable platform. Unfortunately, autonomous lawn mowers still lack intelligent functions known from state-of-the-art research like object recognition, obstacle avoidance, localization, dynamic path planning or speech recognition. In contrast, for cleaning robots there is an active research for vision SLAM [14, 16, 17], however, publications on other topics like visual obstacle avoidance is also scarce.

In order to operate, most autonomous lawn mowers use an area wire emitting an electromagnetic field for boundary definition and for homing to the charging station. Navigation inside the wire is typically random, which is simple and robust but takes longer for a full coverage of the lawn than systematic movement. Obstacles within the mowing area are detected by a collision sensor or just by wheel slip after collisions. These collisions can cause damage to obstacles and lawn as well as scratches on the outer shell of the mower, which leads to an unpleasant image after some time. While this simple obstacle avoidance works, it may require an adaptation of the garden environment to the robot rather than vice versa.

In this work, we investigated the possibility of installing a low-cost camera and processing module on autonomous lawn mowers. The challenge was primarily to achieve robustness for weather-proof robots that can run 24/7, including operation in very diverse environments and difficult outdoor lighting conditions.

2. Related Work

Publications specifically on autonomous lawn mowers are very limited. A major activity in the past was the annual Robotic Lawn Mower Competition that has been hosted by

the Institute of Navigation (ION) from 2004 to 2012. While some of the robots showed good results in mowing coverage and obstacle avoidance, the prototype machines were not built under the consideration of price, weather-proofness or manufacturing constraints. In particular, most participants use costly sensors like D-GPS [21] or laser scanners [4, 3]. In contrast, in this work we target a small, low-cost, low-energy solution which limits the choice of sensors and more importantly also the available computing resources.

Apart from cheap internal sensors like accelerometers, cameras provide rich information about the environment and are cheaper than most active sensors. However, cost limitations lead to low available computing resources which excludes most modern classification algorithms like [2]. Schepelmann et al. [25, 24] use camera images for grass and obstacle segmentation. The authors developed a simple yet efficient grass segmentation algorithm based on color and texture features using image moments. However, the test set was very small and thus robustness and generalization questionable. In [7] color and texture features are fused with data from an infrared camera and a laser for outdoor obstacle detection. The experiments show that the fusion leads to a clear improvement with respect to color and texture features alone, however, the usage of infrared and laser sensors makes the system quite costly.

There is also work using learning techniques to interpret camera data. For example in [13] a classifier is used to segment the camera image into different terrain types. One major problem faced here was the apparent color variation with environmental condition. Thus, they introduced one Gaussian mixture model for each environment state. Unfortunately, this makes the system quite complex and adds another stage that might have a wrong classification. In contrast, LeCun et al. [15] directly trained one convolutional neural network on a robot to reproduce human obstacle avoidance teleoperation. This approach may be more generic, however, the model is more complex and there is no clear quality assessment for different environment conditions, which prevents a quantitative comparison.

In [27, 28] an omni-directional position estimation system for an autonomous lawn mower was introduced. The authors used an off-the-shelf autonomous mower and extended it with an omni-camera and an additional compute unit. In the latest work [28] the equipment was nicely integrated in a near-product fashion. Their work showed good results in position estimation, however, no work on intelligent obstacle avoidance has been done nor did the authors perform a long-term test under different environment and weather conditions.

Another activity is the EU project TrimBot 2020 [1]. The final goal of this ambitious project is to demonstrate a prototype that is able to move in a garden environment, avoiding obstacles and approaching hedges to trim them. Such an approach may lead to more general-purpose service robots, however, high requirements in robustness and safety makes a product in the near future questionable.

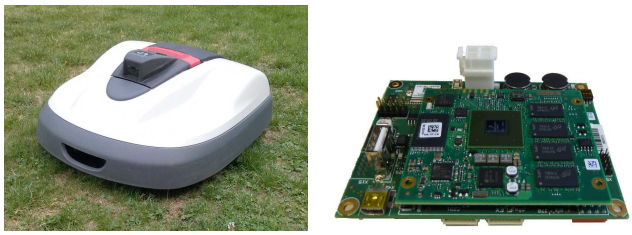

(a) Prototype (b) Compute Board

Figure 1: (a) Prototype autonomous lawn mower with camera module. (b) Internal compute board phyFLEX-i.MX6 on top of a base board providing interfaces.

In contrast to the above mentioned articles, we aim at a low-power low-cost embedded system with special focus on long-term robustness.

3. Approach

3.1. System Overview

For our vision prototype we use an off-the-shelf autonomous lawn mower and extend it with a camera module. Fig. 1a shows our prototype with the camera module attached. The module replaces the maintenance lid and thus seamlessly integrates into the lawn mower body. For computing we use a phyFLEX-i.MX6 (Fig. 1b top board), with an industrial grade Freescale i.MX6 quad-core processor running Linux OS, manufactured by Phytec Messtechnik GmbH (Mainz, Germany). The reason for the industrial grade is the expected high burden on the electronics due to heat and moisture in the garden environment. Especially, in high summer heat can be a major issue as active cooling is not possible because the cut grass snippets would clog the air intake or the fan after a short time.

For seamless integration of the computing board with the autonomous lawn mower, we use a custom base board (Fig. 1b bottom board also by Phytec) that provides a CAN interface to the phyFLEX-i.MX6 board and that allows a direct connection to the battery power line of the autonomous mower. The overall power consumption of the whole module is 5-7W. Hence, the additional load to the mower battery is small and the maximum mowing time is reduced by less than 5%. Furthermore, the direct connection to the mower power line allows a fully autonomous operation of our prototype as the mower recharges autonomously without user interaction.

For image capturing we use a Phytec camera board that features an ON Semiconductors MT9V024 automotive CMOS chip with a resolution of 752x480 pixels and

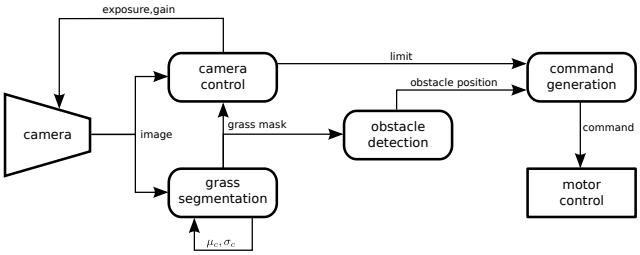

Figure 2: System overview for the visual obstacle avoidance on the autonomous lawn mower.

global shutter together with a Lensagon BM3618C lens (f=3.6mm).

From a software point of view we kept the system as simple as possible in order to get a robust fully autonomous system. Fig. 2 shows a schematic overview of the system's software architecture. There are four main software submodules: camera control, grass segmentation, obstacle detection and command generation. These are explained in more detail below.

3.2. Camera Control

The main challenge for outdoor vision processing are the strongly varying environmental conditions. Firstly, there are potential obstructions on the lens caused by rain drops, moisture, or dirt that disturb the camera image and might lead to misdetections. Although such obstruction can efficiently be detected for example by means of static pixel detection [9], they require a user cleaning action that would make the mower less autonomous.

Secondly, there is the outdoor lighting which on a sunny day in central Europe varies from 100,000lx in the open field to 10,000lx in the shadow, i.e. roughly a difference of three f-stops (relative aperture) in terms of photography. Furthermore, the sun can shine directly into the camera leading to serious sun flares that strongly hamper image algorithms. In order to cope with the difficult outdoor lighting we applied two strategies: scene-dependent camera control and HDR image processing. Note that no white balancing was used since we track the grass hue explicitly (see Sec. 3.3).

The basic idea of the scene-dependent camera control is to set exposure and gain in a way that the relevant scene elements are captured optimally. In our garden scenario this means it is of highest importance to have a good image of the grass so that the grass segmentation algorithm works optimally. All the other scene elements can be under- or overexposed as they will be marked as non-grass anyway. Hence, we employ a closed-loop camera control. As shown in the system graph in Fig. 2, the grass segmentation mask is input into the camera control module. The control module will then use this mask to compute the average brightness of all grass pixels and calculate new exposure and gain values in order to reach a defined target brightness value. Such masking is often used for camera control and calibration in omni-cameras [22, 20] which contain black image parts where the omni-mirror does not project any light onto. In the next time step the grass segmentation will get images with the new exposure and gain setting thus closing the camera control loop.

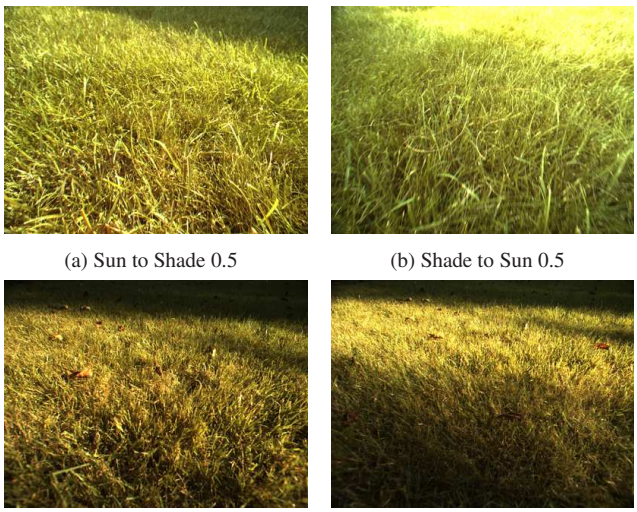

(a) Sun to Shade 0.5 (b) Shade to Sun 0.5

(c) Sun to Shade 0.25 (d) Shade to Sun 0.25

Figure 3: Influence of brightness target setting on image information loss. While a target setting of 0.5 (a), (b) leads to problems for shade to sun transitions (b) due to overexposure, a target setting of 0.25 (c), (d) leads to seemingly darker images. However, as Fig. 4 shows, even the dark areas carry sufficient information for the segmentation.

Another important aspect is the choice of the target brightness value for the camera images, which is often set to 0.5. This value offers good visibility of the camera images on a computer screen for humans. However, computer algorithms do not care for good on-screen visibility. A good target exposure value should have an similar amount of f-stops towards the brightest and darkest pixel values. The darkest reasonable value is determined by image noise, while the brightest values should prevent color distortions caused by overexposure. We experimentally found a target value of 0.25 to give good results. With this setting, pixels can be four times brighter than the target gray value before they get overexposed, and pixels four times darker than the target still have a value well above the chip's noise level.

Fig. 3 depicts images for driving from a shady area to a sunny area and the other way around. Fig. 3a and Fig. 3b show images captured with a target brightness of 0.5 for sun to shade transition and shade to sun transition, respectively. While the sun to shade transition is without problems, in the

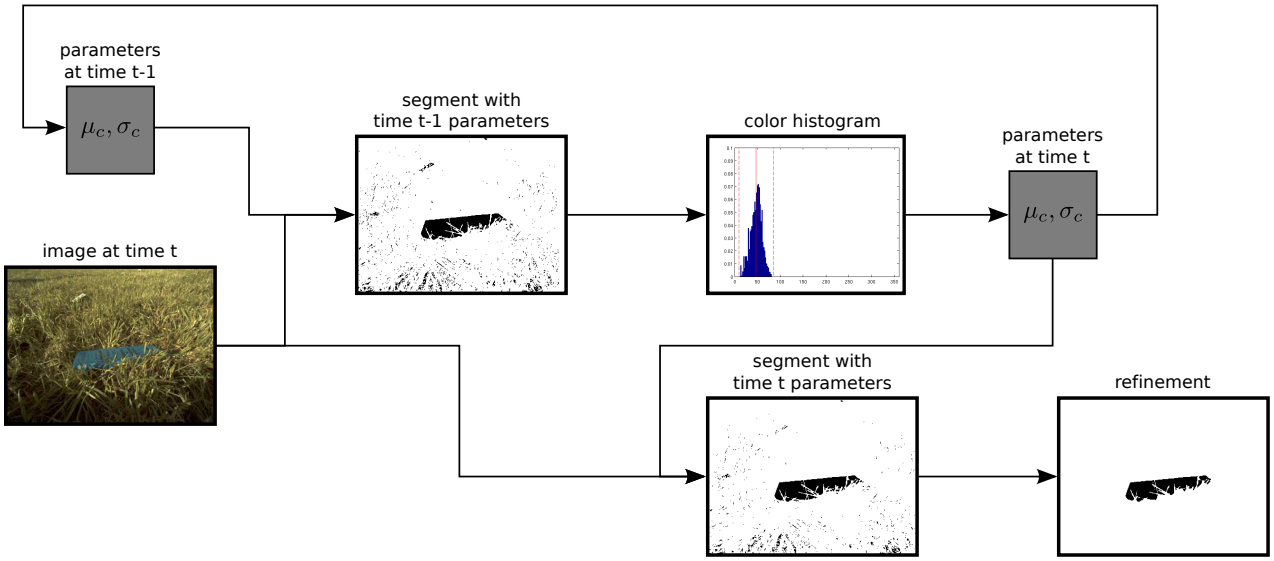

Figure 4: Mean shift grass color tracking and grass segmentation.

shade to sun transition the sunny area gets overexposed and thus the grass cannot be detected correctly anymore. In contrast, using a target brightness of 0.25 (Fig. 3c and Fig. 3d) does not lead to an information loss due to overexposure. The images indeed appear slightly dark but as the example image in Fig. 4 shows, even the dark areas carry sufficient information for the segmentation.

Unfortunately, in some extreme cases on sunny days even a target brightness of 0.25 will lead to overexposure when transiting from the shade to sun area. In order to further improve this situation we propose an HDR control system. Typically, HDR cameras merge information from two consecutively captured images with different exposure setting into one HDR image. The two major problems with this approach are ghost artifacts and more severely an alteration of the color caused by the merging. Since our system relies mainly on color segmentation, the camera's inbuilt HDR mode led to unacceptable results and alternative HDR cameras were not available at comparable cost.

As an alternative, we use a pair of consecutively captured images with distinct exposures and compute the grass segmentation on both of them. Afterwards we merge the two segmentation results to get an HDR result based on the processing and not on the input image. Finally, in rare events when exposure changes are very high, the mower speed is reduced in order to give more time for camera adaptation.

3.3. Grass Segmentation

One key component of our visual obstacle avoidance system is the grass segmentation. We decided to model the grass appearance instead of potential objects/obstacles because the objects have a large variation in appearance. Since grass is typically green and has a high saturation, the hue and saturation components of the HSV space allow easy segmentation, whereas V is less informative because it varies strongly with illumination.

$$M_g(p) = f_h(p) \wedge f_S(p) \tag{1}$$

$$f_H(p) = \begin{cases} 1, & \text{if } h_{min} < H(p) < h_{max} \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

$$f_S(p) = \begin{cases} 1, & \text{if } s_{min} < S(p) \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

This means that a pixel p is marked as grass in the grass mask $M_g(p)$ if the hue H is in a range defined by $[h_{min}, h_{max}]$ and if the saturation S exceeds a threshold s_{min}. For saturation a range is not necessary as grass typically has a very high saturation, i.e. s_{min} is quite large.

However, due to changing lighting conditions and different state of grass health the color and saturation may change during the operation of the autonomous lawn mower. Consequently, fixed segmentation parameters are not feasible. Here we decided to adopt the mean shift algorithm [5] and use it in a temporally continuous update scheme in order to continuously track the color of the lawn. In particular, we track the mean hue μ_h and mean saturation μ_s as well as the variance of the hue σ_h^2 and the variance of the saturation σ_c^2. From these we derive the segmentation parameters as follows:

$$h_{min} = \mu_h - 3 * \sigma_h \tag{4}$$

$$h_{max} = \mu_h + 3 * \sigma_h \tag{5}$$

$$s_{min} = \mu_s - 3 * \sigma_s \,. \tag{6}$$

Figure 5: Example grass segmentation results (white: grass, black: non-grass/obstacle).

Fig. 4 shows an example adaptation loop of the color tracking. At time t a new image is captured and the color parameters from the last time step $t-1$ are used to get an initial segmentation of the scene into grass (white) and non-grass (black). From the grass pixels a new distribution is drawn from which current color parameters are computed. These new parameters are then used to get a better segmentation of the current image into grass and non-grass. Due to image noise the segmentation will be noisy. We solve this problem with a simple connected component analysis [26]. All regions smaller than a threshold will be assigned to the opposite group, i.e. grass becomes non-grass and vice versa. In all our experiments we used a threshold of 50 pixels which roughly corresponds to an obstacle size of 5cm (or 30cm^2 effective area).

In this work, we use a Gaussian model for the grass color distribution. Thus, the mean and variance of the color is computed by:

$$\mu_c^t = \sum_{p \in M_g^t} C^t(p) \tag{7}$$

$$\sigma^{2}{}_c^t = \sum_{p \in M_g^t} (\mu_c^t - C^t(p))^2 \,, \tag{8}$$

where $c \in \{h, s\}$ and accordingly $C \in \{H, S\}$. Fig. 5 shows some example segmentation results of our approach.

One big advantage of using the mean shift algorithm for tracking the grass color parameters is processing speed. The actual segmentation involves only the conversion of the RGB image into HSV space and some threshold operations using the estimated parameters. The update of the parameters can also be calculated very efficiently as it mainly consists of addition and multiplication operations and only a few divisions. This makes the whole segmentation fast and suitable also for very low-cost embedded processors that do not have a floating point unit. On our embedded board the computation without hardware acceleration takes roughly 10ms on 376x240 pixels using one ARM core (800MHz).

Figure 6: From left to right: corridor mask, left mask, right mask. The corridor mask is used to determine which obstacle pixels fall into the path of the autonomous mower. The left and right masks are used to determine whether the obstacle is more on the left or right side so that the mower turns the other direction.

3.4. Obstacle Avoidance

Given the segmentation of an input image into grass and non-grass areas, the system has to derive driving commands for the autonomous mower control for visual obstacle avoidance. First, it is necessary to interpret the segmentation mask. Then the camera module needs to communicate with the motor controller to inform it about obstacles. We decided to keep the interaction simple by creating a contact-free vision avoidance behavior similar to the standard avoidance behavior after collision detection because this has proven to be very robust. This means that the camera module just triggers the random turn maneuver like a collision would. However, in contrast to the collision sensor, the camera module can easily extract some information about the position of the obstacles allowing obstacles avoidance before collision. In order to keep up with the fast processing speed of the segmentation we also decided for a simple and fast solution.

We use three fixed masks: a left mask, a right mask and a corridor mask (see Fig. 6). The corridor mask corresponds to the future driving path of the autonomous lawn mower on a plane. It depends on the width of the mower and the position of the camera. Actually, one would need a 3D information to do this correctly, but we found that a simple ground plane assumption is often sufficient because the driving speed and accordingly the brake distance of autonomous lawn mowers is not very high. For detecting an obstacle we just multiply the corridor mask with the inverse grass mask. Note that this method may fail when objects are not connected to the ground (e.g. an overhanging branch) or when the slope in front of the robot changes extremely.

Additionally, the ground plane assumption allows us to assign a rough distance value to the obstacle by the y-position of the lowest obstacle pixel in the multiplied grass and corridor mask. The distance estimation allows a more smooth visual avoidance by reducing speed during obstacle approach and by turning just before the object.

We use the left and right mask in order to give the motor controller a hint concerning a favorable turn direction. If more obstacle pixels are on the right side the mower should

Figure 7: Camera example views from small scale test.

rather turn left and vice versa.

The system saves logs of internal states and events for later analysis. Saving image data in our long-term test, however, was limited by available mass storage. We thus keep a ring buffer of the last few seconds and only save this data when an event occurs (collision or vision avoidance). Nevertheless, the logs accumulated to a total of 1TB of storage and required automated statistical analysis. We found it helpful to also automatically generate HTML reports of the logs to manually inspect the context of single events, e.g., when annotating obstacle categories and false positive events.

In order to circumvent complex priority negotiations between the camera module and the mower's motor controller we switch off the vision system at the boundary wire and during homing while driving along the boundary wire. We also excluded night time operation in our experiments. Nevertheless, preliminary tests showed successful operation with an additional headlamp.

4. Experiments

4.1. Small Scale Test

In a first and relatively small test, we evaluated the obstacle avoidance performance of the camera systems. This test took place in eight private gardens in the Frankfurt area but also included additional objects placed in the gardens for more efficient testing (see Fig. 7).

The overall detection rate of the camera system was 88% with 11.9 false positive detections per hour. Regarding some especially relevant categories, 99% of bare hand and feet were detected by the camera system, and standing humans were detected with 86%.

4.2. European Field Test

The above test confirmed the working principle of the approach but did not guarantee sufficient performance in long-term unsupervised operation with naïve testers in uncontrolled gardens. Thus, we designed a long-term test over one season in eight European countries (see Fig. 8 and Fig. 9) with more than 20 robot units with camera system in as many different gardens. We evaluated performance, robustness, and checked for systematically unmowed areas. For evaluation of the testers subjective opinions, each tester

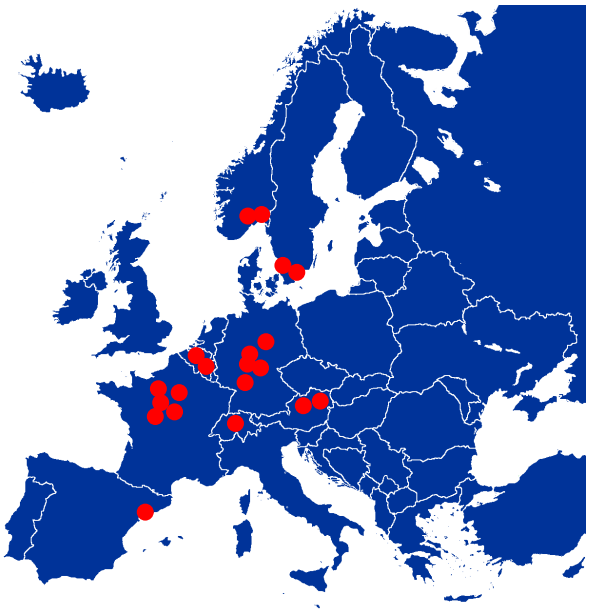

Figure 8: Map of test site locations

Figure 9: This figure shows some sample views of the European test gardens.

was given a questionnaire.

The total recorded operation time was 3,494 hours. This test differed from the earlier one: gardens, seasons, and weather conditions were more diverse, fewer obstacles were excluded by area wire (especially trees), and the robots were not maintained by developers. On average, the robot mower without a vision system would have 0.8 collisions per minute. With visual obstacle avoidance, the collisions were reduced by 76.9%. Obstacle categories, avoidance rates, and frequencies are listed in Table 1. Small objects like toys and tools are recognized with highest accuracy. These objects are often not recognized by a collision sensor, thus possibly causing damage to the blades and objects. People and animals are also easily detected, which was noted by testers as giving a more friendly image. In contrast, green and brown objects, like plants and trees, are on the lower end of reliably detectable objects with our

Table 1: Obstacle occurrences and avoidance rates. Data based on manually annotated samples from each test site.

object class	avoidance rate	occurrence
toys, tools	96.3%	4.0%
people, animals	95.6%	1.3%
pots	92.8%	1.0%
furniture	90.7%	9.8%
misc objects	86.9%	37.5%
poles	75.4%	10.5%
plants (w/o trees)	60.0%	8.5%
trees	59.9%	27.4%
all	**76.9%**	100.0%

color-based system.

4.2.1 Encountered Problems and Non-Problems

During our large scale field test we encountered some unexpected problems, while some expected issues did not occur. For example, no test unit had problems with water proofness or overheating.

On well-kept and sufficiently irrigated lawns without leafs, we measured only 0.15 false positive avoidance events per minute, which were mostly caused by glares and shadows. However, many of the gardens were in poorer condition (earthy and sandy patches without grass) and some were at times almost completely covered in leafs. Hot and dry weather, especially in France, caused completely yellow-brown grass color in some lawns for multiple weeks. On average, including these real-world conditions, the false positive avoidance rate was 0.57 events/minute, dominated by grassless patches and those with dry brown grass (52.0%). Leafs (21.8%), glares (12.3%), and shadows (11.7%) caused less false avoidances and were evaluated as less critical, since their position is usually not fixed and a mulching robot mower typically has to visit the same spot multiple times for good results. In 1.3% of the cases we found strong slope changes causing false avoidances and surprisingly only 0.05% were caused by wet or dirty lens conditions or an insect on the lens. Arguably, the mower should detect very adverse conditions and stop or reduce mowing.

Overall, we found that false positive avoidances caused only a minor loss in mowing time. In no case did we find, or get report of, unmowed garden areas. This is in line with our observation that most false positives are caused by shadow borders, which are not stationary. However, as expected, gardens with distinct lawn areas separated by walkways or driveways (but within one area wire) would not be mowed completely. In these cases, we recommended to program the mower's starting point within each distinct area on different weekdays (a standard feature of the mower by following the area wire for a certain distance before starting to mow). Possible future extensions here would be a classification of walkways as driveable or location-based exceptions for visual obstacle avoidance. All units showed cutting result similar to the autonomous mower without camera, which is an important requirement for using a camera on autonomous lawn mowers because the cutting performance should not be reduced. However, there is one aspect that needs to be tackled in the future. Usually autonomous mowers leave some uncut area around obstacles because of the safety distance between blades and outer robot body. We found that we set our avoidance distance too conservative leading to a larger uncut region around obstacles. This is in particular true for obstacles the autonomous mower without camera would not detect with the collision sensor, e.g. small toys (see Fig. 10). However, this image demonstrates a reliable and high detection performance, since otherwise such areas would not exist.

Figure 10: Uncut area around a toy left on the lawn. On the one hand side, this shows that we chose the avoidance distance too conservative. On the other hand, it shows a reliable obstacle detection.

One case that our system currently does not cover is obstacle avoidance while turning. Since we use a forward looking camera, the system cannot see obstacles to its side, which can lead to a turning into an obstacle. In the future, we want to tackle this problem by implementing an obstacle memory for storing obstacles just passed by.

Another issue, which was expected, is avoidance of green or brownish obstacles with a color similar to the garden's grass. The majority of misdetections occurred with bushes and trees. Note that trees are often covered with lichens or moss near the ground, which shifts the bark color towards brown and green. In the small scale test, we found that these cases are easily avoided with a stereo camera system using the algorithm in [8]. Alternatively, texture or object recognition could reduce this problem.

4.2.2 Smartphone

We additionally developed a smartphone app for interacting with the mower. The app shows a live video of the camera and an optional color-coded overlay of the classification of "mowability". This app also allows a user to switch to remote control (RC) mode. In order to avoid problems with network configurations, we provided an Android smartphone with preinstalled apps to each tester. Usage statistics showed that more than 90% of the testers used remote control of the autonomous lawn mower very rarely and only early in the testing period. Questionnaire results also confirm that RC is perceived as not important to most users, and that most users prefer a completely autonomous operation. However, we found that during development, in early testing, and for remote support the smartphone app was more efficient and easier to handle than interaction with remote login over ssh with a laptop.

4.2.3 Tester Feedback

Most testers returned questionnaires after testing and gave additional oral feedback. Most notably, the majority of testers assessed the quality of the mower with camera system as significantly higher than without and would purchase such a product if available. A minority, however, expected near-perfect performance ("If he has eyes, why does he not see this?") and would prefer a stereo system with higher performance even with increased cost. Most testers state that the autonomous lawn mower and people (adults and/or children) are on the lawn during operation. While safety of the mower is guaranteed by the collision sensors, the perceived friendliness increases with the camera system. Interestingly, many testers gave a pet name to their mower and many noted that the camera module increases the perception of the robot as pet-like and that it increased their emotional attachment. We conjecture that two factor play a role: the increased complexity of the behavior, and the association of the camera with an eye.

5. Conclusion

In this work we have presented an embedded camera module prototype for autonomous lawn mowers for visual obstacle avoidance. The major challenges were to make the system robust against the varying outdoor weather and lighting conditions. To tackle these we devised a closed-loop between the camera control and a color-based grass segmentation algorithm. The camera control optimizes the visibility of the grass in the scene by taking the grass segmentation output into account. In addition the grass segmentation uses a mean shift algorithm to track the changes in grass color due to the changing environment conditions. The system was implemented on a small module which was attached to an off-the-shelf autonomous lawn mower. The module is connected via CAN for communication with the mower control ECU and the module is connected to the mower power for utilizing the self-charging capabilities making the whole system able to run 24/7.

In order to test the developed system we first did a small scale test that we also used to parameterize the algorithms and to finalize the prototype. Afterwards we conducted a large field test by deploying 20 units to eight European countries. These prototypes where installed for one season and recorded all avoidance events from both collision and vision sensors. The analysis of the data revealed that we could reduce collisions by 77% which strongly increases comfort (less noise from collisions) and product appearance (less scratches on the outer mower body).

Camera systems in outdoor environments face a multitude of challenges. However, they provide a rich source of environmental information allowing many functions, e.g. obstacle avoidance. Our long-term field test demonstrated that image processing can be sufficiently robust for an outdoor service robot. Compared to many modern computer vision approaches, the color-based approach requires drastically less computing power, which makes it realistic for low-cost low-power systems. An extension to texture and object recognition is interesting for the future. However, we found that the color-based approach has a robustness advantage when the lens is wet or dirty. In these rare cases, the image becomes blurred, which may lead to a complete breakdown of performance for some algorithms, whereas the color-based approach showed graceful degradation of performance.

Future steps encompass a local map to memorize obstacles for improved obstacle avoidance in blind spots, e.g. during turning maneuvers, and robust self-localization [18, 19]. We also want to further investigate the feasibility to use a stereo camera for tackling the issue of detecting obstacles close to grass color and to improve the obstacle localization with respect to the mower. Finally, we also want to enable the mower to recognize certain types of obstacles for a class specific behavior. For example to mow over weed while avoiding the users favorite flower.

Acknowledgments

We would like to thank H. Shimamura and M. Yamamura from Honda R&D Japan for providing ECU software modifications to the autonomous lawn mower. Furthermore, we want to thank Honda Motor Europe for providing test machines. In addition, we would like to thank J. Barter from Honda Motor Europe for supporting us in contacting people for using their gardens as test environment. Finally, we want to thank T. Wakitani, R. Borsdorf and E. Wrase from Honda R&D Europe for supporting the project with their domain expertise.

References

[1] http://trimbot2020.webhosting.rug.nl/. 2

[2] V. Badrinarayanan, A. Kendall, and R. Cipolla. SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2017. 2

[3] S. Baichbal. Mapping algorithm for autonomous navigation of lawn mower using Sick laser. Master's thesis, Wright State University, 2012. 2

[4] J. A. Beno. CWRU Cutter: Design and control of an autonomous lawn mowing robot. Master's thesis, Case Western Reserve University, 2010. 2

[5] D. Comaniciu and P. Meer. Mean shift analysis and applications. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1197–1203, 1999. 4

[6] H. Şahin and L. Güvenç. Household robotics: Autonomous devices for vacuuming and lawn mowing. *IEEE Control Systems*, 27(2):20–96, 2007. 1

[7] C. S. Dima, N. Vandapel, and M. Hebert. Classifier fusion for outdoor obstacle detection. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 665–671, 2004. 2

[8] N. Einecke and J. Eggert. A multi block-matching approach for stereo. In *Intelligent Vehicles Symposium (IV)*, pages 585–592, 2015. 7

[9] N. Einecke, H. Gandhi, and J. Deigmöller. Detection of camera artifacts from camera images. In *Intelligent Transportation Systems (ITSC)*, pages 603–610, 2014. 3

[10] J. Fink, V. Bauwens, F. Kaplan, and P. Dillenbourg. Living with a vacuum cleaning robot. *International Journal of Social Robotics*, 5(3):389–408, 2013. 1

[11] M. Hägele, editor. *World Robotics 2016 Service Robots*. International Federation of Robotics Statistical Department, 2016. 1

[12] R. W. Hicks and E. L. Hall. Survey of robot lawn mowers. In *Intelligent Robots and Computer Vision XIX: Algorithms, Techniques, and Active Vision*, 2000. 1

[13] P. Jansen, W. van der Mark, J. C. van den Heuvel, and F. C. A. Groen. Colour based off-road environment and terrain type classification. In *IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 216–221, 2005. 2

[14] M.-J. Jung, H. Myung, S.-G. Hong, D.-R. Park, H.-K. Lee, and S. Bang. Structured light 2D range finder for simultaneous localization and map-building (SLAM) in home environments. In *Micro-Nanomechatronics and Human Science, and The Fourth Symposium Micro-Nanomechatronics for Information-Based Society*, pages 371–376, 2004. 1

[15] Y. LeCun, U. Muller, J. Ben, E. Cosatto, and B. Flepp. Off-road obstacle avoidance through end-to-end learning. In *NIPS*, pages 739–746, 2005. 2

[16] H.-K. Lee, K. Choi, J. Park, and H. Myung. Self-calibration of gyro using monocular SLAM for an indoor mobile robot. 1

[17] S. Lee, S. Lee, and S. Baek. Vision-based kidnap recovery with SLAM for home cleaning robots. *Journal of Intelligent & Robotic Systems*, 67(1):7–24, 2012. 1

[18] B. Metka, M. Franzius, and U. Bauer-Wersing. Outdoor self-localization of a mobile robot using slow feature analysis. In *International Conference on Neural Information Processing (ICONIP)*, pages 249–256, 2013. 8

[19] B. Metka, M. Franzius, and U. Bauer-Wersing. Improving robustness of slow feature analysis based localization using loop closure events. In *International Conference on Artificial Neural Networks (ICANN)*, pages 489–496, 2016. 8

[20] A. J. Neves, A. Trifan, and B. Cunha. Self-calibration of colormetric parameters in vision systems for autonomous soccer robots. In *Robot Soccer World Cup*, pages 183–194, 2013. 3

[21] G. Newstadt, K. Green, D. Anderson, M. Lang, Y. Morton, and J. McCollum. Miami RedBlade III: A GPS-aided autonomous lawnmower. *Positioning*, 8(1):25–34, June 2010. 2

[22] N. Nourani-Vatani and J. Roberts. Automatic camera exposure control. In *Australasian Conference on Robotics and Automation (ARAA)*, pages 1–6, 2007. 3

[23] E. Prassler, M. E. Munich, P. Pirjanian, and K. Kosuge. Domestic robotics. In *Springer Handbook of Robotics*, pages 1729–1758. 2016. 1

[24] A. Schepelmann. Identification & segmentation of lawn grass based on color & visual texture classifiers. Master's thesis, Case Western Reserve University, 2010. 2

[25] A. Schepelmann, R. E. Hudson, F. L. Merat, and R. D. Quinn. Visual segmentation of lawn grass for a mobile robotic lawnmower. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 734–739, 2010. 2

[26] L. D. Stefano and A. Bulgarelli. A simple and efficient connected components labeling algorithm. In *International Conference on Image Analysis and Processing*, pages 322–327, 1999. 5

[27] J. Yang, S.-J. Chung, S. Hutchinson, D. Johnson, and M. Kise. Vision-based localization and mapping for an autonomous mower. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3655–3662, 2013. 2

[28] J. Yang, S.-J. Chung, S. Hutchinson, D. Johnson, and M. Kise. Omnidirectional-vision-based estimation for containment detection of a robotic mower. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6344–6351, 2015. 2