Road Following with Blind Crawling Robot

Martin Stejskal   Jakub Mrva   Jan Faigl

***Abstract*— In this paper, we consider road following to autonomously navigate a mobile robot through an environment while keeping the robot on the specified road. Contrary to existing approaches based on a forward looking camera, we consider the problem for a technically blind walking robot without any exteroceptive sensors. The only feedback considered is an estimation of tactile information that is determined from the robot servo drives. The proposed control law is based on an on-line classification of the previously learned terrains which is utilized to identify a situation when a robot starts to crawl off the desired road terrain. The controller steers the robot to keep its body and all its legs on the road while crawling forward with a constant velocity. The experimental results support feasibility of the proposed minimalistic approach and allows the robot to autonomously navigate in an outdoor environment and follow urban park pathways and avoid off-road parts.**

***Index Terms*— terrain classification, crawling robot**

I. Introduction

Road following can be considered as one of the first desired tasks of mobile robots towards fully autonomous vehicles. First results on this problem can be dated to 1985 [1] and since that, many approaches have been proposed to keep the forward moving robot on the road. The problem depends on a definition of the road, since the road can be of various types of surface a robot can traverse [2]. For vision-based techniques, which can be considered as a vast majority of the existing road following approaches, the road is usually considered as a surface of different color or texture for which reactive based approaches [3], adaptive techniques [4], and on-line learning procedures [5] have been proposed.

Contrary to vision-based techniques, we consider a minimalistic approach for a technically blind mobile robot without sensors. The proposed approach is based on the tactile information provided by the robot actuators themselves. In particular, we consider a hexapod walking robot, see Fig. 1, equipped with intelligent servo drives that are able to provide estimation of the current joint torque, which is the only feedback used in the proposed autonomous road following.

Similarly to vision-based techniques with pre–learned terrain classes [6], [7], also the proposed approach relies on the previously learned classes. In particular, an extension of [8] to rough terrains [9] enabled by the adaptive motion gait [10] is considered. On top of this combination, a new reactive-based controller is designed to steer the robot position to be on the road surface whenever the robot detects its left or right legs start to crawl off the specified terrain.

Authors are with Department of Computer Science, FEE, Czech Technical University in Prague, CZ `stejsma6 | jakub.mrva | faiglj@fel.cvut.cz`

The presented work has been supported by the Czech Science Foundation (GAČR) under research project No. 15-09600Y.

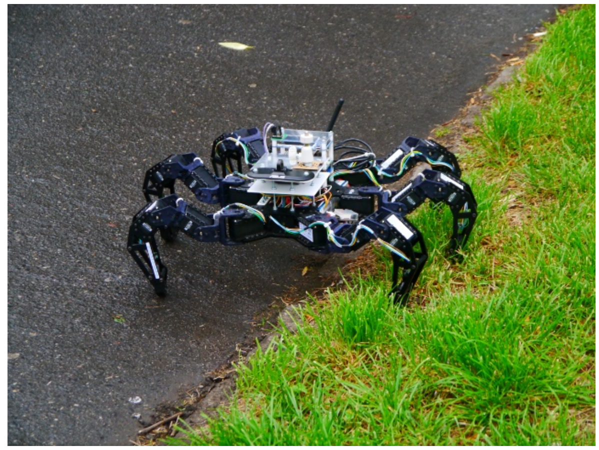

Fig. 1. Hexapod walking robot during autonomous road following

The combination of the proposed controller with the on-line terrain classification based on servo drives feedback only is considered as the main contribution of this paper. The approach allows the robot to autonomously navigate on pathways without any exteroceptive sensors. The proposed system has been experimentally verified in a laboratory environment and further validated in real outdoor environments, where the robot was able to follow pathways in an urban park without completely leaving the pathway. Even though the proposed approach will unlikely be used as a standalone solution, it provides a minimalistic approach that can be accompanied by additional, e.g., vision-based, techniques to improve its performance by complementary methods.

The paper is organized as follows. An overview of related work is provided in the next section. The problem statement with a brief description of the platform is in Section III. An overview of the adaptive motion gait and an extension of the on-line terrain classification for rough terrains are presented in Section IV. The proposed road following controller is in Section V. Experimental results are reported in Section VI. Concluding remarks are in Section VII.

II. Related Work

Although road following is widely studied approach in the context of autonomous cars, there is not too much work on road following with a blind robot, especially for multi-leg crawling robots. Therefore, an overview of the most related work on the terrain classification focused on proprioceptive sensing and crawling robots is presented in this section.

In [7], authors discuss various sensor modalities in the terrain classification and they reported the best performance is provided by gyros. Besides, they also report that inertial sensors allow to distinguish a terrain type if the terrain cause distinct vibrations. Authors of [11] also report that a terrain classification can be based on the vibrations. In both these approaches, a wheel mobile robot has been utilized, albeit vibration measurements are also considered with legged robots without adaptive gaits [12].

For crawling robots, the important part of road following (with blind robot) is the ability to traverse different uneven terrains. Although several approaches for dealing with a rough terrain have been proposed, e.g., see overview in [13], the main point of our interest are approaches based on proprioceptive sensors that do not rely on range measurements provided by the robot exteroreceptors. Here, the key to provide a stable gait is to detect the leg contact with the terrain and avoid high torques at the joints.

An on-line force based foothold adaptation has been utilized in [14] to provide a smooth contact of the leg with the foothold. Authors of [15] propose to avoid usage of direct force sensors by additional passive actuators that are utilized for measuring the ground reaction force. The authors consider hexapod robot with 18 controllable degrees of freedom accompanied with passive actuators in each leg. Then, the motion of each leg is driven by a force threshold-based position controller with the ground reaction force estimated from the passive actuator.

In a very similar way, even more minimalistic approach has been proposed in [10], where instead of additional sensors or actuators, the robot active actuators are directly employed in robot motion and the torque for detection of the leg contact point with the foothold is estimated.

Tactile sensors provide a feedback about the robot interaction with the terrain and thus, they have been used for terrain classification. A single vibrating leg detached from the robot body with force sensors attached at the leg tip has been considered with the motor current at the knee joint in [16]. A 6-DOF torque-force sensor was used in [17] for a discriminant analysis between six types of terrain. However, these approaches also utilize additional sensors, which do not fit with our intention of the minimalistic approach.

A different and minimalistic approach that utilizes only proprioceptive sensors built within the actuators was proposed in [8]. The authors show that no additional sensors are needed for successful terrain classification on flat terrains, which was verified also on rough terrains by authors of [9].

From the reactive control perspective, energy consumption depends on the terrain being traversed as well as on the gait being used and can thus be used as a cue for changing control strategy (switching the gait) [18]. The authors, however, showed results only for simple non-challenging terrains and used gaits designed for perfectly flat surfaces.

In this paper, we propose a combination of the terrain classification method [8] with the adaptive motion gait [10] to address the road following for a blind hexapod walking robot. Both these approaches are based solely on proprioceptive sensing using only feedback from the robot active actuators. The classification is utilized in the designed decision-making strategy for steering the motion to keep the robot on the road using only the information about the classified terrain.

III. Problem Statement

The addressed problem of the road following is considered in the context of autonomous navigation of a blind hexapod walking robot crawling on various surfaces. The robot does not have any exteroceptors nor information about its heading. The only feedback available are measurements provided by the active robot actuators. The idea of the proposed approach is to use the feedback to identify the terrain the robot is currently traversing and thus, allow a terrain classification. However, a reliable terrain classification using proprioceptive sensors needs data from crawling on a new terrain. Therefore, it is not expected the robot will stay entirely on the road.

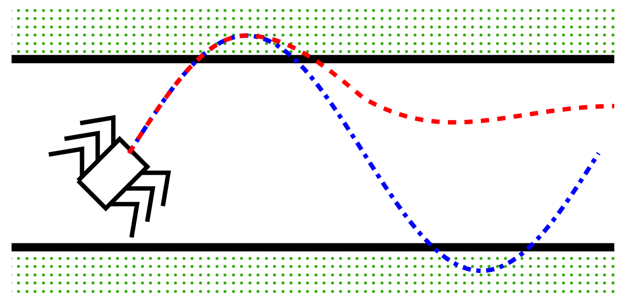

Fig. 2. Possible behaviour of a blind robot after sensing an off-road terrain. the blue dash-dotted trajectory denotes a simple reactive strategy for which the robot keeps crawling an opposite direction than the off road. The red dashed line shows a behaviour of the desired control strategy.

The utilized hexapod robot has six legs symmetrically attached on the left and right sides of its body. Thus, the robot can utilize the actuator feedback from left and right legs which allows to identify the border of two terrains [9]. Therefore, we can distinguish four basic classes characterizing the robot state according to the road following:

- *On road* – the robot has legs on the desired road surface;
- *Off road* – the robot has legs off the road;
- *Left off road* – the left legs are off the road;
- *Right off road* – the right legs are off the road;

For brevity, the terrain classes are denoted as $\mathcal{T}_\mathcal{C}$=$\{\mathtt{On}, \mathtt{Off}, \mathtt{Off}_{\mathtt{left}}, \mathtt{Off}_{\mathtt{right}}\}$. Notice, multiple surfaces can be represented by a single terrain class from $\mathcal{T}_\mathcal{C}$, e.g. dirt or grass can both represent `Off` class.

We consider a robot moving with a constant forward velocity, and the problem is to design a control strategy based on the terrain classes such that the robot will follow the road. Since there is no other information about the heading available, the proposed approach can utilize only the previous estimation of the traversed terrain. Two possible strategies are shown in Fig. 2. Without considering the information about previously traversed terrains, the robot can be simply steered to the opposite direction than the border was detected, as it is shown by the blue dash-dotted curve in Fig. 2. However, our goal is to achieve a behaviour of the robot that will look like the red dashed curve in Fig. 2.

The proposed approach is based on the utilized adaptive motion control [10] combined with the on-line terrain classification [8] that has been extended for crawling on rough terrains [9]. An overview of these methods is presented in the next section.

IV. Hexapod Robot, Adaptive Motion Gait and On-Line Terrain Classification

The proposed road following method is considered with a low-cost walking robot visualized in Fig. 1 that is based on the PhantomX Hexapod Mark II platform. The robot consists of six legs, each with three joints realized from the Dynamixel AX-12A servos. Each servo drive contains a built-in position controller and it communicates over the half-duplex serial bus. Beside the position control itself, the servos are able to measure the position error, current speed and torque. This feature provides information about the leg motion without a need of any additional sensors.

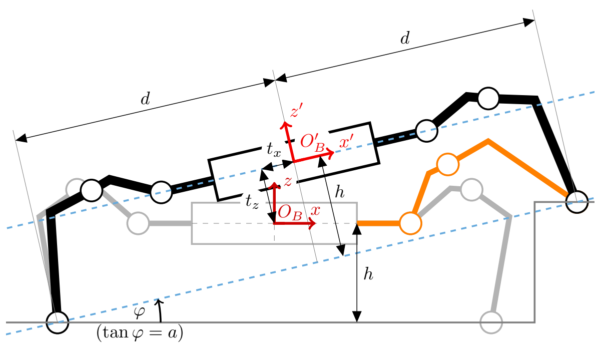

Fig. 3. Simplified schema of body leveling. When a leg (right gray) reaches a new foothold (orange), the body posture is adjusted by changing the configuration of the legs (black) while keeping the same distances h and d.

To allow the robot to traverse various terrains, the available information about the servo position error is utilized in the adaptive gait [10]. This gait is based on a regular tripod gait with a modified phase of the leg motion towards the ground in which the contact between the leg and terrain surface is used to stop the motion and avoid high torque values at the joints. The gait works as follows. First, the legs are moving up and then forward. Then, the legs are moving down towards the ground. During this phase, the position error is read from the servos. The ground detection is based on monitoring the magnitude of the position error and the ground is detected for its predefined threshold. After all legs stand on the footholds, the final phase is activated to level the robot body to fit with the actual positions of the legs. The body leveling is schematically visualized in Fig. 3.

The main advantage of the considered adaptive gait [10] over a regular gait is its ability to traverse any terrain with a constant speed. The robot can pass obstacles of size that respects its physical capabilities. Besides, the motion of the robot is much smoother with the adaptive gait when walking on uneven terrains.

A. Terrain Classification with Adaptive Motion Gait

The position error is used also in the proposed on-line terrain classification that is based on the approach proposed by the authors of [8], which has been combined with the adaptive motion gait [10] in [9]. Therefore, only a brief overview of the terrain classification and used features are provided in this section to make the paper more self-contained.

In [8], small errors in the position control of all servo drives of the front legs are measured in the time domain to get interpolated samples at the frequency of 100 Hz. The samples are then used to construct a feature vector for the classification. Features are computed as statistics of the data samples that correspond to a particular part of the gait cycle that is divided into 16 equally wide segments. Respective segments from the last three gait cycles are joined together and basic statistics of all data samples that fall within are computed yielding in 5 values (features) for each segment (i.e., the minimum, maximum, mean, median, and standard deviation). This is done for each servo which gives 480 gait-phases features, i.e., 2 (legs) $\times$ 3 (servos per leg) $\times$ 5 (features) $\times$ 16 (segments).

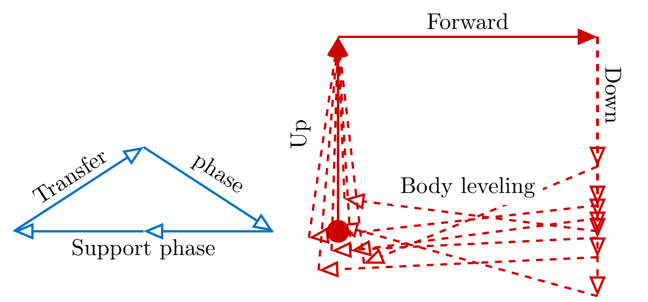

(a) Default gait (b) Adaptive gait

Fig. 4. Comparison of the leg trajectory using a regular default gait and an adaptive gait

Beside these features, the authors of [8] also consider additional features found in the frequency domain enabled by a regular gait. However, the adaptive motion gait does not provide a periodic behaviour, because the required time to finish the full gait cycle depends on the terrain the robot is crawling, see Fig. 4. Due to the variances of the gait phases, which depends on the roughness of the terrain the robot is traversing, we cannot rely on a uniform partitioning of the gait phases into 16 segments for the feature extraction as in [8]. On the other hand, we can utilize the gait phases of the adaptive gait that consists of 4 phases per each leg triplet, and the data from one gait cycle can be therefore divided into 8 segments according to the gait phases [9].

Having the feature vectors from several trials of the robot crawling on different terrains, the Support Vector Machine (SVM) algorithm is utilized to train the models of the terrain classification. The models can be then utilized in on-line terrain classification by feeding them with the current feature vector. Notice, due to the way how the feature vector is constructed, the robot needs to finish full gait cycle to make a decision about the terrain it is crawling.

B. On-line Terrain Classification for Road Following

The problem of the road following has been firstly tackled using a terrain classification separately for left and right legs of the robot. In this setup, models were created for each robot side. However, early results on this separate classification with easily distinguishable terrains were not promising and the attitude proved to be very unreliable. This classification fails because the walking robot is one interconnected body and any change of the terrain on one side affects also the leg trajectory (and its measurement) on the other side.

Therefore, a more suitable approach is to consider feature vectors from legs of both robot sides in a single terrain model and a border between two terrains have to be trained by the model as a regular terrain. This results in training of four classes $\mathcal{T}_\mathcal{C}$={`On`, `Off`, $\texttt{Off}_\texttt{l}$, $\texttt{Off}_\texttt{r}$}.

Regarding the influence of the gait and robot motion to the terrain classification, it is worth mentioning that the proposed approach is based on training of the classifier using data from a straight walk of the robot. Although this simplify acquisition of the training datasets, it may results in a less accurate terrain classification during steering the robot back to the road, especially at the border of two terrains. An experimental validation of this training and terrain classification is reported in Section VI.

V. Proposed Road Following Control Strategy

The proposed decision-making strategy is based only on the information about the type of the terrain the robot is traversing. Our goal is to control the robot in the way of the red dashed curve depicted in Fig. 2.

The robot forward velocity is considered constant, and therefore, the problem of road following is to determine the robot angular velocity that will steer the robot to keep all its legs on the road. Since we do not have any prediction of the road curvature, we propose to estimate a relative direction based on the last N environment observations provided by the terrain classification. Hence, individual observations can contribute to the estimation of the required steering direction and we propose the control law as

$$\dot{\theta} = k_p \sum_{i=1}^{N} \frac{1}{\eta i} s_i w_i, \tag{1}$$

where $\dot{\theta}$ is the robot angular velocity directly applicable to the gait controller, k_p is the proportional control gain, η is an aging factor, s_i is a sign of the steering weight w_i. The aging factor η is applied to incrementally decrease influence of the previously detected terrains to the final steering action. The steering weight w_i and its sign s_i are determined according to the current and previously classified terrains T_i, $1 \le i \le N$, and they are determined as follows.

After each gait cycle, the robot determines a new terrain type $T_i \in \mathcal{T}_\mathcal{C}$. For a new terrain type T_i, the current steering weight w_i is set to one of three possible values:

$$w_i(T_i) = \begin{cases} W_{off} & \text{if } T_i \text{ is } \texttt{Off}; \\ W_b & \text{if } T_i \text{ is } \texttt{Off}_\texttt{left} \text{ or } \texttt{Off}_\texttt{right}; \\ w_{on}(i) & \text{if } T_i \text{ is } \texttt{On}; \end{cases} \tag{2}$$

where constants W_{off} and W_b stand for the off road and border weight, respectively. The weight W_{off} represents a sharp turn, while the weight W_b represents a less aggressive correction on the road border, therefore $0 \le W_b \le W_{off}$.

The weight $w_{on}(i)$ depends on the previously traversed terrains and its determination follows the idea that once a robot is on the road for a couple of full gait cycles, it should keep its heading. On the other hand, if the robot was previously off the road it should keep the turning radius to return to the road. Finally, for the robot mostly on the border, its direction should only slightly change. We propose to determine $w_{on}(i)$ as follows

$$w_{on}(i) = \begin{cases} W_{off} & \text{if } n_{off} \ge N_{off}; \\ W_b & \text{if } n_b \ge N_b \text{ and } n_{off} < N_{off}; \\ 0 & \text{otherwise -- \textit{keep the direction}}; \end{cases} \tag{3}$$

where n_{off} and n_b are the number of times the terrain has been determined as `Off` and $\texttt{Off}_\texttt{left}$, $\texttt{Off}_\texttt{right}$ respectively for $1 \le i \le N-1$. The thresholds N_{off} and N_b denote the numbers of terrain samples T_i for which the history of the terrain classification is considered mostly off road and mostly as border terrain, respectively.

Finally, for each weight its sign $s_i \in \{-1, 1\}$ is determined based on the desired robot orientation. The positive sign is for clockwise, while the negative sign stands for the counter-clockwise direction. The sign is based on the current terrain type T_i and the previously determined type T_{i-1}:

$$s_i = \begin{cases} -1 & \text{if } T_i = \texttt{Off}_\texttt{left} \text{ -- \textit{left side border}}; \\ 1 & \text{if } T_i = \texttt{Off}_\texttt{right} \text{ -- \textit{right side border}}; \\ -s_{i-1} & \text{if } T_i = \texttt{On} \text{ and } T_{i-1} \ne \texttt{On}; \\ s_{i-1} & \text{otherwise}. \end{cases} \tag{4}$$

If a transition from the border or the off-road terrain to the on road is captured, the sign is set to the opposite of the previous one. Also regarding the rule (3), a new action then acts against the previous one. This causes the robot to get partially back to the road border to compensate the whole maneuver. Once $w_i(T_i)$ and s_i are established, the control action is computed using (1).

A. Limits of the Proposed Control Strategy

The proposed control strategy covers most of the situations of crawling on various terrains during road following. The only situation impossible to handle is a direct transition from `On` to `Off`, where no information about the side is available.

Regarding the terrain classification, it is worth mentioning that the SVM model is trained only from straight walks. Therefore when an angular velocity exceed some limit, a confusion of the terrain classification can occur. It is caused by a different trajectory of the leg with respect to the one during a straight walk used in terrain learning. Considering this, the control action cannot be applied continuously. It is necessary to alter between terrain classification and heading correction for a single robot step. Each of these steps takes exactly one gait cycle. However, for traversing the `On` terrain and $w_{on} = 0$, only the terrain classification is performed in a step that allows to use a more frequent information about the terrain type in the classification and thus, the robot reacts more promptly to changes of the terrain.

VI. Experimental Results

Two different scenarios have been considered to evaluate the proposed steering control to avoid off-road terrains and following the road. At first, we considered easily distinguishable terrains that were artificially prepared in a laboratory. After that, we have considered outdoor experiments in an urban park where several trials have been performed. In both scenarios, an experiment was considered as a successful if the robot smoothly return back to the road after leaving it.

The SVM model for the terrain classification was trained off-line from the previously collected data. However, during the road following, the classification and steering action are employed in on–line data processing. Only two front legs are used for the classification, which is performed at the end of the gait cycle utilizing data from the last three cycles.

All the algorithms have been implemented in C++ and executed on the robot onboard ARM-based computer Odroid U3 with 1.7GHz Quad-Core processor and 2GB of RAM.

Prior the evaluation, we experimentally tuned the parameters of the proposed control strategy. The used weights and thresholds values of (1) and (3) are as follows: the length of the terrain history $N = 10$, the proportional control gain has been set to $k_p = 0.5$, the aging factor to $\eta = 2$, the weights to $W_{off} = 1$, $W_b = 0.5$, and $N_{off} = 5$, $N_b = 4$.

A. Laboratory Experiment

A verification of the on-line classification and control strategy has been evaluated on the terrains that can be easily recognized. Using surfaces easy to distinguish minimizes the influence of the misclassified samples on the decision-making strategy. Such conditions were provided by a wooden board and pillows placed on the floor, see Fig. 5. The pillows are very soft in contrast to the wood, and therefore, both surfaces are easy to recognize.

Fig. 5. Indoor testing environment – wood and pillows. According to $\mathcal{T}_C$ classes, the surfaces were assigned as follows: `On`– wood, $\texttt{Off}_{\texttt{right}}$– wood–pillows, $\texttt{Off}_{\texttt{left}}$– pillows–wood, `Off`– pillows.

Repeated walks over all four learned terrain classes from $\mathcal{T}_C$ approved the expectation of easily distinguishable terrains. The robot has been left to walk the border between `On` and `Off` in both directions and the terrain classification has been monitored. The accuracy of the terrain classification in this case reached 100%, which validates a feasibility of the proposed control strategy. It also allowed fine tuning of the parameters prior the deployment in an outdoor environment.

TABLE I

Confusion matrix of 2–fold cross-validation with overall accuracy 96.2%. Class names are shortened in the table in the following way: D – Dirt, A – Asphalt, G – Grass

Class	G	D – A	D	A – D	A	G – A	A – G
G	99	0	0	0	0	0	0
D – A	0	116	7	3	0	1	0
D	0	9	62	0	0	0	0
A – D	0	4	1	82	0	0	1
A	0	0	0	0	112	0	0
G – A	0	0	0	0	0	117	0
A – G	0	0	0	2	0	0	122

B. Outdoor environment

The outdoor experiments have been conducted in an urban park with asphalt pathways surrounded by a grass or a dirt. An asphalt has been chosen to represent the `On` class and a dirt and grass have been marked as the `Off` class. The terrains are visualized in Fig. 6 with their assignment to $\mathcal{T}_C$.

The fidelity of the model trained in this manner has been firstly evaluated by the two-fold cross validation. The overall achieved accuracy is 96.2%. As it can be seen from the confusion matrix in Table I, the most challenging terrain for the classification is a dirt. Almost 86% of misclassified samples was related to dirt. Regarding the assumptions described in Section IV, the terrain classification can be considered as more difficult since the transition from asphalt to dirt was usually flat.

Other false predictions in the confusion matrix in Table I belong to the same $\mathcal{T}_C$ class, thus they are acceptable. Such confusions arise from the similarity of the borders between asphalt and other off-road terrains.

During the outdoor experiments, the robot has been driven to the borders similar to those depicted in Fig. 6. Due to significant differences between the terrains, like in Fig. 6a, borders are easy to deal with. Also the borders between asphalt and grass terrains (shown in Fig. 6b and Fig. 6c) are classified reliably. The most challenging terrain to classify is the dirt. During the experiments, we included basically two types of the border asphalt–dirt. The first has been a strict border with a small hump as shown in Fig. 6d and it has been easy to deal with. The second type of the asphalt–dirt border is shown in Fig. 6e. This border is flat, blurred and there is a dirt spilled over the edge and the transition of the robot between the terrains is much slower.

In addition, we also evaluated various angles how the robot approached to the border for each type of the border. Due to the length of the robot step and ability to classify the terrain from the last three whole gait cycles, for angles larger than approximately 45°, the algorithm loses the information about the side, where the off road was, because the border is passed too fast, and a straight transition from the on road to the off road is reported. Therefore, it is impossible for the robot to return to the road in the way described in Section V.

(a) Border stone (b) Flat border with grass (c) Steeper border with grass (d) Strict border with dirt (e) Blurred border with dirt

Fig. 6. Borders of the terrains considered in the outdoor experiments. Assignment of the surfaces to $\mathcal{T}_\mathcal{C}$ classes was done as following: $\mathtt{On}$ – asphalt, $\mathtt{Off}_{\mathtt{right}}$– asphalt–grass and asphalt–dirt, $\mathtt{Off}_{\mathtt{left}}$– grass–asphalt and dirt–asphalt, $\mathtt{Off}$– grass and dirt.

The reported results from the performed evaluation and achieved results are further documented in the video attached to the paper. The result from one of the experiment captured in the accompanied video is shown in Fig. 7. The robot approached the grass, $\mathtt{Off}_{\mathtt{left}}$ class has been detected and the control strategy pushed the robot back to the road. After the asphalt (as the $\mathtt{On}$ class) has been detected again, the robot aligned back approximately to the road direction.

Fig. 7. A trail of the robot path following an asphalt pathway. The traversed trajectory is highlighted by the red line.

VII. Conclusion

In this paper, we propose a simple road following method for a technically blind crawling robot. Contrary to a majority of already existing methods we consider only proprioceptive sensing; in fact, only the feedback provided by the active servo drives is utilized. The proposed minimalistic approach utilizes the feedback to detect a contact point of the leg with the foothold and also for the terrain classification of the previously learned terrains. The proposed decision-making strategy computes the steering action from the history of the classified terrains. The experimental results from the outdoor environments support feasibility of the proposed approach and the control strategy keeps the robot on the specified road and also provides a rough estimation of the road direction captured in the recorded history of the terrain types.

Even though the proposed method allows the blind crawling robot to autonomously navigate through an outdoor environment while keeping on the road, the method can be rather considered as a complementary approach to other road following techniques to improve their robustness based on the robot experience with negotiating particular terrains. Therefore, further methods based on exteroceptive sensing are subject of our future work.

References

[1] R. Wallace, A. Stentz, C. E. Thorpe, H. Maravec, W. Whittaker, and T. Kanade, "First Results in Robot Road-Following." in *IJCAI*, 1985, pp. 1089–1095.

[2] M. Ososinski and F. Labrosse, "Automatic Driving on Ill-defined Roads: An Adaptive, Shape-constrained, Color-based Method: Automatic Driving on Ill-defined Roads," *Journal of Field Robotics*, vol. 32, no. 4, pp. 504–533, 2015.

[3] T. Krajnik, J. Blazicek, and J. Santos, "Visual road following using intrinsic images," in *ECMR*, 2015, pp. 1–6.

[4] D. Lieb, A. Lookingbill, and S. Thrun, "Adaptive Road Following using Self-Supervised Learning and Reverse Optical Flow." in *Robotics: Science and Systems*, 2005, pp. 273–280.

[5] W. Mou and A. Kleiner, "Online learning terrain classification for adaptive velocity control," in *IEEE International Workshop on Safety Security and Rescue Robotics (SSRR)*, 2010, pp. 1–7.

[6] R. Hadsell, P. Sermanet, J. Ben, A. Erkan, M. Scoffier, K. Kavukcuoglu, U. Muller, and Y. LeCun, "Learning long-range vision for autonomous off-road driving," *Journal of Field Robotics*, vol. 26, no. 2, pp. 120–144, 2009.

[7] L. Ojeda, J. Borenstein, G. Witus, and R. Karlsen, "Terrain characterization and classification with a mobile robot," *Journal of Field Robotics*, vol. 23, no. 2, pp. 103–122, 2006.

[8] G. Best, P. Moghadam, N. Kottege, and L. Kleeman, "Terrain classification using a hexapod robot," in *Proceedings of the Australasian Conference on Robotics and Automation (ACRA)*, 2013.

[9] J. Mrva and J. Faigl, "Feature extraction for terrain classification with crawling robots," in *Proceedings ITAT 2015: Information Technologies - Applications and Theory*, 2015, pp. 179–185.

[10] ——, "Tactile sensing with servo drives feedback only for blind hexapod walking robot," in *RoMoCo*, 2015, pp. 240–245.

[11] C. C. Ward and K. Iagnemma, "Speed-independent vibration-based terrain classification for passenger vehicles," *Vehicle System Dynamics*, vol. 47, no. 9, pp. 1095–1113, 2009.

[12] F. Garcia Bermudez, R. Julian, D. Haldane, P. Abbeel, and R. Fearing, "Performance analysis and terrain classification for a legged robot over rough terrain," in *IROS*, 2012, pp. 513–519.

[13] D. Belter, P. Labecki, and P. Skrzypczynski, "Adaptive Motion Planning for Autonomous Rough Terrain Traversal with a Walking Robot," *Journal of Field Robotics*, 2015, DOI: 10.1002/rob.21610.

[14] A. Winkler, I. Havoutis, S. Bazeille, J. Ortiz, M. Focchi, R. Dillmann, D. Caldwell, and C. Semini, "Path planning with force-based foothold adaptation and virtual model control for torque controlled quadruped robots," in *ICRA*, 2014, pp. 6476–6482.

[15] L. Palmer and M. Palankar, "Blind hexapod walking over uneven terrain using only local feedback," in *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2011, pp. 1603–1608.

[16] M. A. Hoepflinger, C. D. Remy, M. Hutter, L. Spinello, and R. Siegwart, "Haptic terrain classification for legged robots," in *ICRA*, 2010, pp. 2828–2833.

[17] K. Walas, "Tactile sensing for ground classification," *Journal of Automation, Mobile Robotics & Intelligent Systems*, vol. 7, no. 2, pp. 18–23, 2013.

[18] N. Kottege, C. Parkinson, P. Moghadam, A. Elfes, and S. P. Singh, "Energetics-informed hexapod gait transitions across terrains," in *ICRA*, 2015, pp. 5140–5147.