

SVMs for Vibration-based Terrain Classification

Christian Weiss, Matthias Stark, and Andreas Zell

University of Tübingen, Department of Computer Science, Sand 1, 72076 Tübingen

Abstract. When an outdoor mobile robot traverses different types of ground surfaces, different types of vibrations are induced in the body of the robot. These vibrations can be used to learn a discrimination between different surfaces and to classify the current terrain. Recently, we presented a method that uses Support Vector Machines for classification, and we showed results on data collected with a hand-pulled cart. In this paper, we show that our approach also works well on an outdoor robot. Furthermore, we more closely investigate in which direction the vibration should be measured. Finally, we present a simple but effective method to improve the classification by combining measurements taken in multiple directions.

1 Introduction

In outdoor environments, a mobile robot should be able to adapt its driving style to the ground surface. Some surfaces like asphalt are flat and not slippery, and thus the robot can safely traverse them at high speeds. Other surfaces, like sand or gravel, are dangerous because they are slippery and/or bumpy. To prevent accidents or damages to the hardware, the robot has to traverse these surfaces carefully at low speed. Such hazards that originate from the ground surface itself can be called *non-geometric hazards* [1]. A system which can determine the type of the current or forthcoming ground surface therefore greatly contributes to the safety of a robot.

Iagnemma and Dubowsky first suggested to detect non-geometric hazards by vibrations that are induced in the robot while traversing the terrain [2]. The vibrations are different for different terrain types, and characteristic vibration signals can be learned for each terrain type. Based on the learned model, the terrain class of a newly collected vibration signal is estimated. Such a method can be used as a stand-alone classifier or to supplement other sensors.

Commonly, accelerometers are used to measure vibrations. The accelerometers can be placed at the wheels or the axes of the robot, as well as on the robot's body. Mostly, the acceleration is measured in up-down (z) direction. However, this paper shows that it could be better to use the acceleration measured in front-back (x) or left-right (y) direction.

Brooks and Iagnemma presented vibration-based terrain classification for low-speed planetary rovers [3]. They use *Principal Component Analysis* (PCA) to reduce the dimensionality of the vibration data and *Linear Discriminant Analysis* (LDA) for classification. Sadhukhan and Moore proposed an approach

that is based on *Probabilistic Neural Networks* [4,5]. They used an RWI ATRV-JR robot driving up to 0.8 m/s. In [6], we presented an approach that uses *Support Vector Machines* (SVM) for classification. Stavens *et al.* suggested a method for vehicles driving up to 35 mph [7]. However, they focus on determining the roughness of the terrain to adapt the speed of the vehicle, instead of grouping the terrain into classes.

In [6], we presented experiments on data collected by a hand-pulled cart. This cart has relatively hard rubber wheels which lead to clear vibration signals. The big, air-filled wheels of common outdoor robots, however, are likely to dampen the vibrations. Therefore, this paper presents experiments that use our method on vibration data collected by an RWI ATRV-JR outdoor robot. We also more closely investigate the influence of different robot velocities on the classification. Additionally, we evaluate in which direction the acceleration signals should be measured. Finally, we propose a simple method to improve classification if the robot is able to measure accelerations in multiple directions.

2 Terrain Classification Method

This section summarizes our terrain classification approach presented in [6] and suggests a simple way to improve classification by using multiple directions of vibrations.

2.1 Basic Method

Our terrain classification approach has two phases: training and classification. Training is computationally intensive and therefore an offline step. Classification is very fast and can be done online. In the training phase, we learn characteristic vibration signals for known terrain types. For this purpose, the robot traverses different surfaces and collects vibration signals. To collect the vibrations, we use an accelerometer that works at 100 Hz. In the next step, we split the acceleration signals into segments, where each segment corresponds to 1 s of robot travel. In our case, this leads to 1×100 -sized vectors. We then label each vector with the terrain type it corresponds to. Fig. 1 shows example acceleration vectors for some terrain types. Except for grass, they appear very similar to a human.

Next, we transform the raw acceleration signals to the frequency domain. In [6], we compared different transformations. Despite the fact that a 128-point Fast Fourier Transform (FFT) led to worse results than other transformations on the cart data, we found the FFT to work best for the ATRV-JR data. After applying the FFT to each vector, we normalize each feature (= frequency component) to mean 0 and standard deviation 1. The normalization prevents features with high magnitude from dominating the training.

Next, we train a Support Vector Machine (SVM) [8,9] on the feature vectors. As kernel function we use a Radial Basis Function (RBF) $k(x, y) = \exp(-\|x - y\|^2/2\sigma^2)$, where x and y are two feature vectors. We tune the width σ of the RBF kernel together with the soft margin parameter C (which regularizes

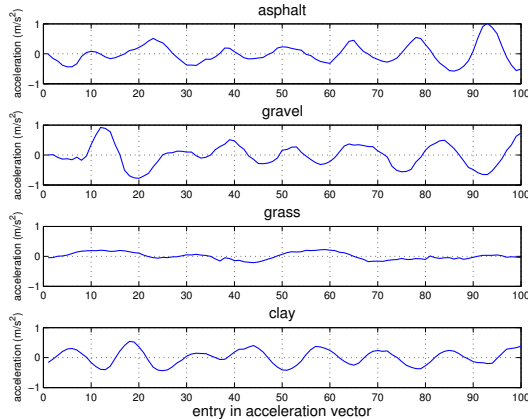


Fig. 1. Some example acceleration vectors for asphalt, gravel, grass and clay.

the trade-off between maximizing the margin between two classes and minimizing the training error) of the SVM by a grid search. The grid is defined by $\log_2 \sigma \in \{\hat{\sigma}/4, \dots, 4\hat{\sigma}\}$ and $\log_2 C \in \{-2, \dots, 14\}$, where $\hat{\sigma}$ is set such that $\exp(-D/2\hat{\sigma}^2) = 0.1$. D denotes the length of the feature vectors. Each candidate parameter vector (σ, C) on the grid is evaluated by 5-fold cross-validation. The grid parameters are standard ones that are often used in other applications, too. As SVM implementation, we use LIBSVM [10].

In the classification phase, the robot traverses unknown terrain and collects vibration signals. Once per second, it creates a 1×100 test vector from the acceleration values taken during the last second and transforms the vector using the FFT. Additionally, the robot normalizes the feature vector using the same parameters used during training. Then, the trained SVM classifies the test vector and returns the estimated terrain type.

2.2 Combining Different Directions of Vibrations

In vibration-based terrain classification methods, vibration is commonly represented by the acceleration measured in up-down (z) direction. The reason is that bumps in the terrain are likely to have their major effect in up-down direction. However, the accelerations measured in other directions, e.g. front-back (x) or left-right (y), can also be used to capture the vibration. Our experiments presented in Section 3 show that these accelerations may even be more suitable for classification than the data measured in z direction.

In addition, many acceleration sensors are able to measure accelerations along three axes simultaneously. For robots equipped with such a sensor, we propose a simple but effective method to improve classification. For each terrain segment, we collect the acceleration signals along all three axes, in our case front-back (x), left-right (y), and up-down (z). Then, we transform the signals individually by




Fig. 2. a) Our RWI ATRV-JR outdoor robot “Arthur”. b) The terrain types we used in our experiments: 1) indoor floor, 2) asphalt, 3) gravel, 4) grass, 5) paving, 6) clay.

the FFT. Next, we concatenate the transformed feature vectors and normalize the features. We then train the SVM on these feature vectors.

In the classification phase, we perform the same steps to get a test vector containing information about all three directions of accelerations. Finally, we classify the test vector using the SVM.

3 Experimental Results

To get experimental data, we used our RWI ATRV-JR outdoor robot (Fig. 2 a). We mounted an Xsens MTi sensor on an aluminium plate on top of the robot. The sensor measures the acceleration in x , y and z direction simultaneously at 100 Hz.

In total, we collected 10225 terrain vectors, some of them in mid-July and some in the beginning of December. Each terrain vector corresponds to 1 s of robot travel. The terrain vectors differ in the type of terrain and in the velocity of the robot. As terrain types, we used indoor floor, asphalt, gravel, grass, paving, clay (a boule court), and the situation in which the robot did not move (Fig. 2 b). The velocity of the robot was one of 0.2 m/s, 0.4 m/s or 0.6 m/s. Tab. 1 shows an overview of the dataset.

For evaluation, we used 10-fold cross validation, i.e. for each experiment, we split the data into 10 parts and evaluated 10 sub-experiments (= folds). In each fold, we used 9 parts for training and the 10th part for testing. The final result is the mean over the results of the individual folds.

Fig. 3 a) and Tab. 2 show the results of a first set of experiments, for which we used the three terrain classes grass, clay and gravel. We expected that higher speeds of the robot would lead to stronger vibrations and therefore to clearer vibration signals. However, the 3-class experiments did not confirm this expectation. Data collected at 0.2 and 0.6 m/s could be classified similarly well and

Table 1. Number of samples per class in our dataset

class	0.2 m/s	0.4 m/s	0.6 m/s	total
indoor floor	282	549	581	1412
asphalt	499	513	600	1612
gravel	311	323	392	1026
grass	482	572	631	1685
paving	314	573	567	1454
clay	423	579	605	1607
no motion	199	615	615	1429
total	2510	3724	3991	10225

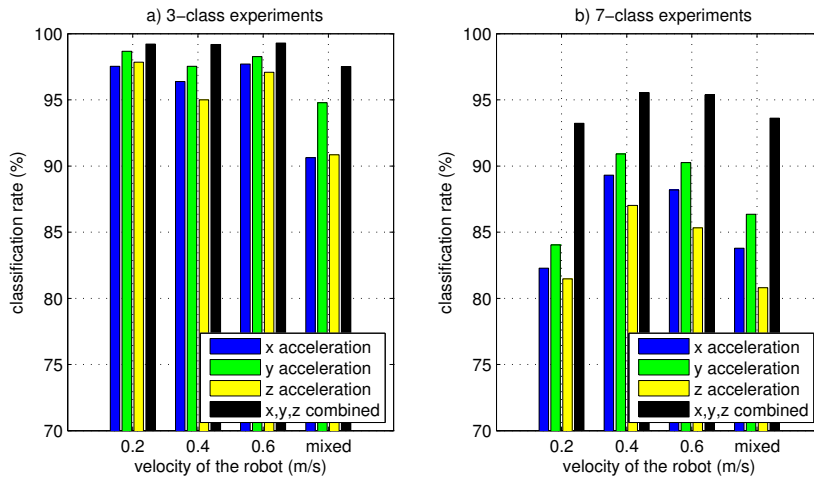


Fig. 3. Experimental results using a) 3 terrain classes, b) 7 terrain classes.

data collected at 0.4 m/s only slightly worse. However, in experiments on all velocities mixed together in one dataset, the classification rates decrease.

The experiments show that the vibration measured in y direction leads to better results than the vibration measured in x or z direction. The y data leads to classification rates of about 98.5% for robot speeds of 0.2 and 0.6 m/s, to 97.54% for 0.4 m/s and to 94.78% on mixed velocities. When combining the data from all three directions, the results get even better. For all individual velocities, the classification rate is higher than 99%. On the mixed dataset, 97.52% of the test vectors are classified correctly.

Fig. 3 b) and Tab. 3 present the results of experiments involving all seven terrain classes. In these experiments, the results for data collected at 0.2 m/s are significantly worse than for data collected at 0.4 or 0.6 m/s. Again, vibration data measured in y direction are classified better than data of the other directions. Additionally, x data yields to better results than z data. The classification rates for the y data are 84.04% at 0.2 m/s, 90.91% at 0.4 m/s and 90.26% at

Table 2. Classification rates (%) of the 3-class experiments

measurement direction	0.2 m/s	0.4 m/s	0.6 m/s	mixed
x	97.53	96.39	97.71	90.63
y	98.67	97.54	98.26	94.78
z	97.84	95.00	97.09	90.84
x, y, z combined	99.20	99.18	99.30	97.52

Table 3. Classification rates (%) of the 7-class experiments

measurement direction	0.2 m/s	0.4 m/s	0.6 m/s	mixed
x	82.28	89.31	88.19	83.78
y	84.04	90.91	90.26	86.40
z	81.47	87.02	85.33	80.80
x, y, z combined	93.22	95.55	95.38	93.61

0.6 m/s. When using combined data from all velocities, the classification rates again increase significantly to between 93.2% and 95.5% for the individual velocities. On the mixed dataset, 93.61% of the test vectors are classified correctly, which is an improvement of over 7% over the y direction alone.

Tab. 4 shows the confusion matrix of the 7-class experiment on mixed velocities and using the combination of the x , y and z measurement directions. An entry (i, j) of the confusion matrix shows how often (in %) test vectors belonging to class i were classified as class j . According to the confusion matrix, the classes “no motion” and grass are unproblematic, because they are confused with other classes in very few cases. Indoor floor is wrongly classified as asphalt for about 3% of the test vectors and vice versa. Gravel and paving seem to be relatively similar. They are confused with each other in about 7-12% of the cases. For clay, there is no clear trend towards misclassifying it as a particular class.

On a 3 GHz Pentium 4 PC with 1 GB of RAM, classifying one terrain vector takes less than 1 ms. The time for training depends on the dataset. For example, training in the 3-class experiment on y data collected at 0.2 m/s takes about 7 min 4 s. This dataset contains 1216 feature vectors. Another example is the 7-class dataset with combined acceleration directions and mixed speeds. This dataset contains 10225 feature vectors, and training takes about 14 h 24 min.

4 Conclusion

This paper showed that our vibration-based terrain classification method presented in [6] works well on a common outdoor robot. Additionally, we presented a technique to improve the classification by using vibrations measured in different directions.

A comparison between data measured in front-back (x), left-right (y) and up-down (z) direction showed that the y data leads to significantly better classification rates than the z data. However, it is not clear if this result is specific for

Table 4. Confusion matrix (%) of the 7-class experiment with mixed velocities and combined measurement directions

	no motion	indoor	asphalt	gravel	grass	paving	clay
no motion	99.79	0	0	0	0.21	0	0
indoor	0	94.05	2.90	0.43	0.14	0.14	2.34
asphalt	0	3.29	93.18	0.37	0.19	1.12	1.86
gravel	0	0	0	84.80	3.12	11.89	0.20
grass	0	0	0	1.18	98.34	0.36	0.19
paving	0	0	0.96	6.88	0.21	90.44	1.51
clay	0	1.43	0.50	0.93	1.00	1.49	94.65

our robot or if this result can be generalized for robots of other types. Nevertheless, the result shows that it is worth trying some other measurement directions before relying on the up-down vibration. Our experiments showed that if the robot is able to measure the acceleration in multiple directions simultaneously, including all available information into the feature vector significantly improves the classification rate.

For future work, we plan online learning, where the new information of the test phase is integrated into the model. The robot should also be able to notice if it traverses some terrain that it has never traversed before.

References

1. Wilcox B. H.: Non-Geometric Hazard Detection for a Mars Microrover. Proc. AIAA Conf. on Intelligent Robots in Field, Factory, Service, Space. Houston, TX, pp. 675 - 684, 1994
2. Iagnemma K., Dubowsky S.: Terrain Estimation for High-Speed Rough-Terrain Autonomous Vehicle Navigation. Proc. SPIE Conf. on Unmanned Ground Vehicle Technology IV, 2002
3. Brooks C. A., Iagnemma K.: Vibration-Based Terrain Classification for Planetary Exploration Rovers. IEEE Trans. on Robotics, 21(6), pp. 1185 - 1191, 2005
4. Sadhukhan D., Moore C.: Online Terrain Estimation Using Internal Sensors. Proc. Florida Conf. on Recent Advances in Robotics, Boca Raton, FL, 2003
5. Sadhukhan D.: Autonomous Ground Vehicle Terrain Classification Using Internal Sensors. Masters Thesis, Dept. Mech. Eng., Florida State University, 2004
6. Weiss C., Fröhlich H., Zell, A.: Vibration-based Terrain Classification Using Support Vector Machines. Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Beijing, China, pp. 4429 - 4434, 2006
7. Stavens D., Hoffmann G., Thrun S.: Online Speed Adaptation Using Supervised Learning for High-Speed, Off-Road Autonomous Driving. Proc. Int. Joint Conf. on Artificial Intelligence, Hyderabad, India, 2007
8. Cortes C., Vapnik V.: Support Vector Networks. Machine Learning, 20, pp. 273 - 297, 1995
9. Schölkopf B., Smola A. J.: Learning with Kernels. MIT Press, Cambridge, MA, 2002
10. Chang C. C., Lin C. J.: LIBSVM: A Library for Support Vector Machines. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>, 2001