

Adaptive Bayesian Filtering for Vibration-based Terrain Classification

Philippe Komma, Christian Weiss and Andreas Zell

Abstract—Outdoor robots are faced with a variety of terrain types each possessing different characteristics. To ensure a safe traversal a robot has to infer the current ground surface from sensor readings. Recent techniques generate a model which predicts the terrain class from single vibration signals disregarding the temporal coherence between consecutive measurements. In this paper, we present a novel approach in which the final classification relies on the analysis of not only one, but several recent observations. Therefore, the probabilistic framework of the Bayes filter is adopted to the problem of terrain classification. We propose an adaptive approach which automatically adjusts its parameters according to the history of observations. To demonstrate the performance of our method we further describe and compare another technique based on temporal coherence. The evaluation using data collected from our RWI ATRV-Jr robot shows that our approach is both reactive and stable enough to detect fast terrain transitions and selective misclassifications.

I. INTRODUCTION

Recently, a growing number of outdoor applications for mobile robots, like rescue missions or agricultural assignments, has emerged. For mobile robots this implies an increasing demand regarding their driving behavior. In outdoor environments, robots are exposed to a variety of different terrain types. To enable a safe traversal of unknown terrain, the robot should adapt its driving style according to the presence of ground surface hazards like slippery or bumpy surfaces. These hazards are known as non-geometric hazards [1].

The characteristics of the ground surface can be inferred by at least two different approaches. One technique involves a direct estimation of terrain parameters like cohesion or slippage without knowing the exact terrain type the robot is driving on. In the second approach, different terrain types are grouped into classes, each representing a ground surface of a certain degree of hazardousness. Using sensor measurements, a model is generated which predicts the present class from the set of available classes.

Most approaches employ vision [2], [3] or lidar sensors [4], [5] for terrain characterization. However, the primary objective of lidar-based methods is to differentiate between the ground surface and obstacles instead of identifying non-geometric hazards.

Other techniques rely on the interaction between the robot and the terrain. Thereby, the robot "senses" internal and external variations like wheel sinkage or wheel slippage. More commonly, algorithms focus on vehicle vibrations as

originally proposed in [6]. They showed that vibration signatures provide enough information to distinguish between different terrain classes. Usually, accelerometers are used to record vibration data during the robot traversal. The sensors can be attached at the wheels, the axes or the body of the robot.

Several researchers considered terrain classification based on vibration data as an instance of a signal processing task. The methods are usually divided into an offline training and an online test or recall phase. In the training phase a model learns to recognize distinct terrain types from labeled vibration patterns. In the recall phase vibration signals of unknown terrain type can be quickly classified using the generated model.

In the terrain classification approach of [7] and [8], the model was represented by probabilistic neural networks. While the former approach focused on terrain classification for electric powered wheelchairs at speeds of 1 and 2 m/s, the latter approach considered both a slowly moving autonomous ground vehicle driving at 0.5 and 1 m/s and an experimental unmanned vehicle at speeds between 5 and 20 mph. Another classification technique was presented by [9]. They used principle component analysis for both feature extraction and dimensionality reduction. In a second step, the PCA transformation coefficients were adopted to establish a manifold curve representing the terrain behavior at varying speeds. In a previous paper, we experimentally investigated the performance of several feature extraction schemes [10]. These techniques were applied to vibration data recorded at driving speeds between 0.2 and 1.0 m/s. The extracted features constituted the inputs for the terrain classification model represented by a support vector machine (SVM). Further research involved the comparison between different classification techniques like the Naïve Bayes or the k -Nearest Neighbor classifier [11]. It turned out that the SVM outperformed all other techniques with respect to classification performance.

In this paper, we propose an extension of the work presented in [10]. To motivate our research, we first note that none of the previously mentioned approaches makes use of temporal coherence. Each terrain classification step only considers the current observation. It is very likely, however, that the robot traverses the same terrain type for several measurements. Hence, not only current sensor readings should influence the classification, but also past ones. The model which considers temporal coherence must have the characteristics to be responsive to terrain changes at high frequency. As we will see later, basic approaches fulfill these requirements

Ph. Komma, Ch. Weiss, and A. Zell are with the Chair of Computer Architecture, Computer Science Department, University of Tübingen, Sand 1, D-72076 Tübingen, Germany {philippe.komma, c.weiss, andreas.zell}@uni-tuebingen.de

only to a limited extent. Instead, we use a Bayes filter, which has been applied in several other domains like object tracking [12] and robot localization [13]. Both problems can be solved using a particle filter which is an instance of the Bayes filter presented in this work. The Bayes filter allows for an effective combination of probabilities for consecutive individual measurements. The combined probability estimates are then used for predicting the current terrain type.

The remainder of this paper is organized as follows: Sect. II focuses on the concept of Bayes filtering which provides the basis of our algorithm. The terrain classification model is introduced in Sect. III. Experimental results are presented and discussed in Sect. IV. Finally, Sect. V gives conclusions.

II. BAYES FILTERING

In our approach, we adopt a Bayes filter which allows for the estimation of a dynamic system's state from noisy observations. This section summarizes the key elements of Bayes filtering. A detailed description is provided in [14]. Using Bayes filters, the state of a dynamic system at a time t is represented by a random variable x_t . Its uncertainty is denoted by a probability distribution over the state space. Given $t + 1$ sensor readings $u_{0:t} = \{u_0, u_1, \dots, u_t\}$, the estimated target distribution is determined by $p(x_t|u_{0:t})$. Applying Bayes filtering, $p(x_t|u_{0:t})$ can be decomposed in the following way:

$$p(x_t|u_{0:t}) = \alpha_t p(u_t|x_t) p(x_t|u_{0:t-1}), \quad (1)$$

where α_t is a normalizing constant. The derivation of (1) is based on the assumption that given the current state, past observations are independent of present ones. $p(u_t|x_t)$ denotes the likelihood function (or measurement probability) and $p(x_t|u_{0:t-1})$ denotes the predictive distribution representing the current state given past observations. Equation (1) quantifies the correction applied to x_t due to current sensor data.

We obtain the predictive distribution by marginalizing over the previous state:

$$p(x_t|u_{0:t-1}) = \int p(x_t|x_{t-1}) p(x_{t-1}|u_{0:t-1}) dx_{t-1}. \quad (2)$$

Here, $p(x_{t-1}|u_{0:t-1})$ is the posterior distribution from the previous time step. The transition probability describing the system dynamics is denoted by $p(x_t|x_{t-1})$. Bayes filters model the dynamic system by a first-order Markov process. That is, the next state x_t only depends on the immediately preceding state x_{t-1} assuming that the information provided by the state x_{t-1} suffices to predict future states without considering earlier observations.

Equations (1) and (2) yield a recursive formulation of the posterior state distribution given previous and current observations. To determine the posterior probability recursively, an initial probability distribution $p(x_0) \equiv p(x_0|u_0)$ has to be defined. $p(x_0)$ is either initialized with prior knowledge about the initial state or uniformly distributed if no prior knowledge exists.

III. TERRAIN CLASSIFICATION MODEL

In this section, we first describe how the ground surface can be predicted from a single observation using SVMs. Then, we give a detailed description of our Bayes filter approach that combines several successive observations into the final prediction. We also describe another prediction technique which considers the history of observations.

A. Basic Terrain Model Composition

The objective of our approach is to establish a model which estimates the terrain type the robot is currently traversing. The estimation is based on acceleration data collected by an accelerometer mounted on the body of our RWI ATRV-Jr outdoor robot. The acceleration data sampled at a frequency of 100 Hz can be regarded as the vibration which the terrain induces to the robot's body.

In the training phase, the model learns the correct assignment of terrain classes, given a set of n consecutive vibration samples. We therefore split the training data into segments consisting of $n = 128$ samples. In this way, each subset corresponds to 1.28 s of robot travel. Then, we label each segment according to the terrain class it belongs to.

To extract characteristic features from the input signal, each vibration segment is transformed from the time into the Fourier domain by applying the Fast Fourier Transform (FFT). With the determined set of FFT coefficients we then calculate the FFT amplitude spectrum resulting in a 64-dimensional output vector. Given the matrix of all transformed vibration segments, where each row represents a single observation, the columns are scaled to have a mean of zero and a standard deviation of 1. The scaled amplitude spectrum vectors represent the final feature vectors and constitute the inputs for the terrain classification model.

The classification model is represented by a support vector machine. As kernel function we chose the radial basis function (RBF). We determined optimal values for the standard deviation σ of the RBF kernel as well as for the soft margin parameter C by a grid search. Each candidate parameter vector on the grid $(\sigma; C)$ was evaluated by 5-fold cross-validation. As SVM implementation we used LIBSVM [15]. In the recall phase, the robot navigates over unknown terrain acquiring vibration data. It predicts the current terrain type using the terrain classification model generated during training. Therefore, 128 samples are recorded over a period of 1.28 s. For two consecutive segments we permit an overlap of 28 samples to achieve a prediction frequency of 1 Hz. We employ the same preprocessing scheme as in the training phase. First, the FFT amplitude spectrum is determined for each vibration segment. Then we scale the resulting feature components using the parameters obtained during training. Using the technique of [16] the application of the final feature vector to the SVM does not only provide a class prediction, but also an approximation of the posterior $p(x_t = i|u_t)$. This probability distribution denotes the probability that a preprocessed vibration segment u_t belongs to terrain class i . $p(x_t|u_t)$ is embedded into the Bayes filtering framework as presented in the following section.

B. Bayes Filtering Applied to Terrain Classification

In our context, the state vector comprises the class number $i = 1, \dots, k$, where k is the number of terrain classes to discriminate. By this coding scheme we obtain a discrete set of k different states describing the dynamic system. The random variable x_t representing the state vector reveals the uncertainty with which the robot navigates on a certain terrain type. Preprocessed vibration data recorded by accelerometer sensors provide the observations.

To apply Bayes filtering to the problem of vibration-based terrain classification, three probability distributions have to be specified: an initial probability distribution $p(x_0)$ which denotes the probability at which the robot resides on a certain terrain type at time $t = 0$, the measurement probability $p(u_t|x_t)$ defining the likelihood that the vibration data measurement u_t can be observed navigating over a certain terrain type x_t , and the state transition probability $p(x_t|x_{t-1})$ denoting the probability that the robot moves from terrain type $x_{t-1} = j$ to terrain type $x_t = i$.

1) *The initial probability distribution* $p(x_0)$: For the definition of $p(x_0)$ we make no assumptions that the robot is placed on a specific terrain type at time $t = 0$. Hence, $p(x_0)$ is assumed to be uniformly distributed.

2) *The measurement probability* $p(u_t|x_t)$: The distribution $p(u_t|x_t)$ can be learned from training examples using parametric or non-parametric density estimators [17]. Note, however, that the extracted feature vector generated from sensor readings has 64 dimensions. This poses a problem since density function estimation of a high-dimensional random variable is a non-trivial task suffering from the curse of dimensionality. For this reason, we adopt another approach to represent the likelihood function. The key idea is to express the measurement probability in terms of the (estimated) posterior probability, $p(x_t = i|u_t)$, provided by machine learning classifiers. Note that in contrast to the probability density function of u_t , estimates for $p(x_t|u_t)$ are provided by certain classifiers like neural networks and support vector machines with only little additional costs.

Applying Bayes' rule to $p(u_t|x_t)$, we have:

$$p(u_t|x_t) = p(x_t|u_t) \frac{p(u_t)}{p(x_t)}. \quad (3)$$

The term on the right hand side of (3) now depends on the classifier posterior probability and the marginal probability of random variables x_t and u_t , respectively. Given no prior knowledge about the marginal probability of a certain terrain instance, we model $p(x_t)$ as a uniform distribution, i.e., it is assigned the value of $c_1 = 1/k$ for all terrain classes. For $p(u_t)$, we consider the complete Bayes filtering formulation that assigns a probability value to a certain terrain class i , given sensor measurements $u_{0:t}$:

$$p(x_t = i|u_{0:t}) = \alpha_t p(x_t = i|u_t) \frac{p(u_t)}{c_1} p_{\text{pr}}(x_t = i), \quad (4)$$

where $p_{\text{pr}}(x_t)$ is the predictive distribution:

$$p_{\text{pr}}(x_t = i) = \sum_j p(x_t = i|x_{t-1} = j) p(x_{t-1} = j|u_{0:t-1}).$$

Note that the integral of (2) has become a sum since we have a discrete set of possible states.

From (4) we see that $p(u_t)$ is constant for all i and can thus be merged with the constant α to give a new normalizing constant. Introducing $\alpha_t^* = \alpha_t \frac{c_2}{c_1}$ with $c_2 = p(u_t)$ yields the final Bayes filter formulation:

$$p(x_t = i|u_{0:t}) = \alpha_t^* p(x_t = i|u_t) p_{\text{pr}}(x_t = i). \quad (5)$$

3) *The state transition probability* $p(x_t|x_{t-1})$: The transition probability $p(x_t = i|x_{t-1} = j)$ describes the probability of moving from state $x_{t-1} = j$ into state $x_t = i$. Given k states, k^2 individual probabilities have to be defined. These values are stored in a square matrix which is denoted as the transition matrix with elements $m_{ij} \equiv p(x_t = i|x_{t-1} = j)$. The matrix diagonal elements m_{ii} represent the probabilities that the system remains in its current state.

Our approach is based on the heuristic that the terrain class most likely does not change from one measurement to the next. This is realized by assigning relatively large values to the diagonal elements of the transition matrix.

Transition matrix elements m_{ij} with $i \neq j$ are derived from the following two heuristics. First, the probability $p(x_t = i|x_{t-1} = j)$ should increase with the probability to confuse class i with class j . The latter probability can be directly derived from the confusion matrix after training the classifier. The confusion matrix is a square matrix with elements $p_{\text{conf}}(i, j)$. Here, $p_{\text{conf}}(i, j)$ represents the probability that at a given time step t the robot is placed on terrain class i , but the terrain class $j \neq i$ was erroneously predicted by the classifier. Second, a transition from state $x_{t-1} = j$ to state $x_t = i$ should be based on the probability that this transition is detected. In our context, the detection of a ground surface change primarily depends on the probability of classifying the terrain type correctly. This probability can be estimated from the the confusion matrix as well. Formally, our estimate for $p(x_t = i|x_{t-1} = j)$, $i \neq j$, becomes:

$$p(x_t = i|x_{t-1} = j) = \mu (p_{\text{conf}}(i, j) + p_{\text{conf}}(i, i)),$$

where μ is a normalizing constant.

In our implementation, we assign a constant value v to each diagonal element of the transition matrix. Given transition matrix column j , the remaining elements are assigned values proportional to $p(x_t = i|x_{t-1} = j)$ such that $\sum_{i, i \neq j} p(x_t = i|x_{t-1} = j) = 1 - v$.

When applying the Bayes filter, two operations increase the run-time complexity: the evaluation of (5) and the determination of the posterior probability using the pairwise coupling method of [16]. Note, however, that the evaluation of (5) only involves an operation count of order $O(k^2)$ for both multiplications and additions. Since k , the number of classes to discriminate, is typically small, real-time requirements are maintained. Further tests revealed that the determination of the posterior probability influences the terrain class prediction complexity only insignificantly. Hence, we obtain a real-time approach allowing for online terrain classification on a real robot.

C. Adaptive Bayes Filtering

The results presented in Sect. IV indicate that choosing a constant value for the parameter v in our Bayes filter approach is unsuitable if the distances between varying terrain types differ significantly. Instead, we propose an adaptive scheme that selects the parameter v dynamically according to the history of observations. The longer it is assumed that the robot navigates on the same terrain, the larger the value assigned to v should be.

Let $w = 1 - v$ be the probability of a system state change, that is, a change in the estimate of the current terrain class. During robot traversal, w is either increased or decreased by a constant factor $\lambda > 1$, depending on the previous system state ($prevState$) from time $t-1$ and the current and previous predictions performed by the SVM classifier, denoted by $currPred$ and $prevPred$. Thereby, we distinguish between the following four cases:

1) $currPred = prevPred$ and $currPred \neq prevState$:

Here, the history of the last two predictions indicates that the robot navigates on a terrain type which differs from the one represented by the Bayes model. Thus, w is enlarged to increase the probability of switching into another system state.

2) $currPred = prevPred$ and $currPred = prevState$:

The terrain class represented by the previous system state equals the classes predicted by the SVM classifier in the last two steps. Hence, it is likely that no terrain transition has occurred and the system resides in the correct state. We therefore reduce the probability of moving into another state by decreasing w .

3) $currPred \neq prevPred$ and $currPred \neq prevState$:

The last two consecutive predictions indicate a transition from one terrain type to another one. Since the system previously resided in a state that differs from the current prediction the probability of a system state change should be increased. This is achieved by increasing w .

4) $currPred \neq prevPred$ and $currPred = prevState$:

The last case is concerned with situations where short-term misclassifications occurred, the system, however, did not move into the false state. Here, it is assumed that the system still resides in the correct state. Hence, we scale down w to decrease the probability of a system state change.

Note that these rules are only heuristics which can influence the classification performance in a negative way if the abovementioned assumptions do not hold. However, the effect of false assumptions is reduced since these rules do not alter the system state directly, but the *probability* that such a transition will occur.

D. A Summed Posterior Probability Approach

Besides our adaptive Bayes filtering approach we also implemented another technique which considers probability estimates of several individual predictions for the final classification. Let x_t be the predicted terrain class at time step t . We consider c observations collected within a fixed period of time t_{n-c+1} to t_n . Using probability outputs of the SVM

model, we obtain an estimate for the posterior distribution $p(x_t|u_t)$. Given posterior probabilities between time step t_{n-c+1} to t_n , a weighted summed posterior probability approach (WSPP) can formally be defined as:

$$\text{predicted class} = \arg \max_i \sum_t f(t_n - t)p(x_t = i|u_t). \quad (6)$$

In other words, the predicted class becomes the one with the largest posterior probability summed over a distinct period of time. Thereby, each probability estimate is weighted according to the recentness of the associated prediction. The weighting accounts for the time when a measurement was acquired (relative to the actual time t_n) and improves the prediction results if the terrain changes frequently. The weighting coefficients are determined using the weighting function f chosen as $f(x) = \exp(-\frac{x}{\tau})$. Here, τ regulates the decline of the function f subject to an increase in x . The greater τ , the greater is the decline in g , the less is the importance of prior observations. In our experiments, $\tau = c/2$ provided the best experimental results.

E. Terrain Classification Model Evaluation

To investigate the reactivity of our methods based on temporal coherence, the evaluation procedure presented in Sect. III-A has to be modified: We assemble consecutive vibration segments representing the same terrain type to give a travel distance of constant length. Then, assembled segments of varying terrain types are grouped together yielding the final test set. In different experiments, we vary the travel distance covered by a robot before it reaches a new terrain class. Given that the modifications are large enough, we can infer the reactivity of our techniques from the resulting classification performance.

IV. EXPERIMENTAL RESULTS

A. Experimental Setup

In our experiments, an Xsens MTi altitude and heading reference system was mounted on an aluminum plate on top of the robot to measure vibration signals in left-right direction at 100 Hz. We collected data in the middle of July and in the beginning of December. During data acquisition, the robot navigated over six different terrains (Fig. 1): indoor PVC floor, asphalt, gravel, grass, paving, and clay (the surface of a boule court). We introduced a seventh artificial terrain type that represents idleness. In the recall phase, this class yields useful information if the robot is not provided with odometry sensors. To not constrain the model to work at a certain driving speed, we varied the speed between 0.2, 0.4, and 0.6 m/s. In total, the dataset consists of 10633 patterns corresponding to approximately 3 hours of robot navigation. The Bayes filter approach and the weighted SPP approach require the definition of a single parameter each. Using the Bayes filter technique, the value v , which corresponds to the diagonal elements of the transition matrix, has to be specified. Assigning a large value to v is equivalent to increasing the probability that the system state does not change from one measurement to the next, that is,



Fig. 1. The terrain types we used in our experiments: 1: indoor floor, 2: asphalt, 3: gravel, 4: grass, 5: paving, 6: boule court.

that two filtered predictions performed at two consecutive time steps yield the same terrain class. c denotes the parameter for the weighted SPP approach representing the number of observations on which each terrain class prediction relies. The larger c , the longer is the history of considered predictions, the larger is the probability that selective misclassifications can be filtered out due to the increased influence of earlier predictions. On the other hand, the approach becomes less reactive to terrain changes as c increases since the enlarged history of considered observations delays the correct state transition to the new terrain class. In our experiments, we chose v from the interval $[0.2; 0.975]$ with step size 0.025 and $c \in \{3; 5; 7; 9\}$. The adaptive Bayes filter approach automatically chooses an appropriate value for v according to the heuristics presented in Sect. III-C. For the parameter λ which regulates the increase and decline of the complementary probability $w = 1 - v$, we experimentally obtained best true positive rates using $\lambda = 1.8$. We further mapped and bounded the values of v to the interval $[1/k; 0.89]$, where k denotes the number of terrain classes to discriminate. The predefined upper and lower bounds, $1/k$ and 0.89, ensure that the probability of moving into another state neither becomes too large nor too small. At time $t = 0$, w was initialized with $w = 1$.

We also varied the distance the robot had to navigate from one terrain class to the next. This distance is denoted as the *travel distance* d (measured in meters) in the following. In each experiment, d was chosen from the set $d \in \{0; 0.6; 1.2; 2.4; 4; 6; 8; 12; 16\}$. The 0 m experiment describes the worst case scenario for approaches based on temporal filtering. Here, single segments of varying terrain classes are concatenated each representing data acquired during 1 s of robot travel. Since the robot speed varies between 0.2, 0.4, and 0.6 m/s, this experiment includes travel distances of 0.2, 0.4, and 0.6 m. Note that according to the confusion matrix, certain terrain transitions are easier to detect than other ones. Hence, the results depend on the order in which assembled terrain segments of varying terrain type are presented. We minimized this effect by randomly permuting this order and averaging the classification results determined after 20 reruns of a particular experiment.

To quantify the performance of the classification results, we applied the evaluation procedure introduced in Sect. III-E. As quality measure we used the true positive rate (TPR). It is the ratio (measured in per cent) between the number of correct predictions for which the predicted class x_t equals the actual class \hat{x}_t and the number of instances contained in the test set.

B. Results and Discussion

Table I shows the classification results for the proposed approaches based on single observations (SO), adaptive Bayes filtering (ABAY), Bayes filtering (BAY) using a constant parameter v for all experiments, and weighted summed posterior probabilities (WSPP). The results were determined by averaging the true positive rate over the complete set of classes. For the Bayes approach and the weighted SPP approach, we chose two conservative parameter settings. The first one, $v = 0.4$ and $c = 3$, yields reasonable classification results in situations of frequent terrain changes. The selection of the second parameter set, $v = 0.85$ and $c = 5$, results in an improved generalization performance when terrain transitions are less likely. Additional prediction results for further settings of parameters v and c are summarized in Table I.

Note that the true positive rate for the single observation based approach differs between varying experiments. This is due to the model evaluation procedure introduced in Sect. III-E. Vibration segments are concatenated to give a travel distance of constant length. If after several concatenation steps the number of remaining segments does not satisfy this criterion these segments are rejected and thus are not considered when determining the generalization error. An increase in classification performance for travel distances 12 and 16 m indicates that those segments have been left out which decreased the true positive rate.

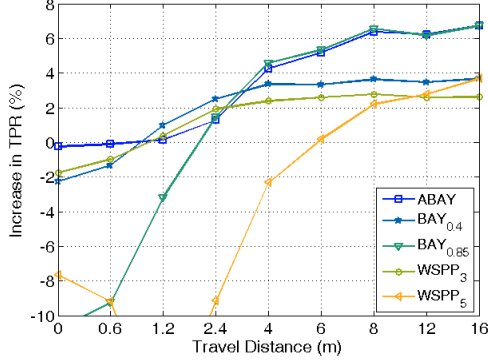
Fig. 2(a) depicts the results when varying the parameters for the proposed approaches. All values denote the increase of the generalization performance in respect of the single observation approach. In the 0 and 0.6 m experiments, the adaptive Bayes approach yields true positive rates similar to those obtained when using the single observation approach. These results indicate that the adaptive Bayes approach chooses small values for the parameter v . Since the probability of changing the system state rises with decreasing v , we obtain a reactive classification system which is able to detect and process high frequent terrain transitions. The assumption that the adaptive Bayes approach prefers small values for v is confirmed by our experiments. They revealed that in $96.0\% \pm 2.0\%$ of all cases, v was assigned the value of $1/k$ in the 0 m experiment. In the 0.6 m experiment, this parameter value was employed less frequently ($78.3\% \pm 1.4\%$), yet more often than the other ones.

Models based on the Bayes approach and the weighted SPP approach are less reactive in situations of high frequent

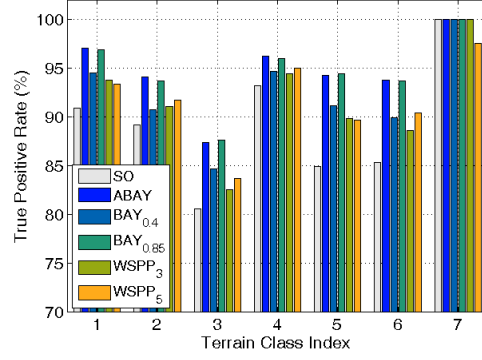
TABLE I

TRUE POSITIVE RATES USING DIFFERENT APPROACHES IN RELATION TO THE TRAVEL DISTANCE (DIST.) AVERAGED OVER ALL CLASSES.

Dist. (m)	0	0.6	1.2	2.4	4	6	8	12	16
SO	86.8	86.9	87.0	87.0	86.9	87.5	86.9	89.1	89.0
ABAY	86.6	86.8	87.1	88.2	90.6	92.0	92.5	94.7	95.0
BAY _{0.4}	84.8	85.7	87.8	89.2	89.8	90.4	90.1	92.2	92.3
BAY _{0.7}	81.1	82.3	86.2	89.1	90.7	91.7	91.9	93.9	94.3
BAY _{0.85}	77.4	78.8	84.2	88.3	90.9	92.2	92.6	94.6	95.0
BAY _{0.9}	75.3	76.9	83.0	87.7	91.0	92.3	93.0	95.0	95.4
WSPP ₃	85.3	86.0	87.3	88.7	89.0	89.8	89.3	91.4	91.4
WSPP ₅	80.2	78.9	72.5	79.1	84.9	87.7	88.8	91.6	92.3
WSPP ₇	74.4	72.7	61.7	72.8	81.3	86.0	87.8	91.0	92.2
WSPP ₉	71.3	69.2	55.3	65.7	75.6	82.0	85.4	89.3	91.1



(a)



(b)

Fig. 2. (a) Comparison of classification results between the proposed methods using varying model parameters related to the single observation approach and (b) true positive rates for individual classes using the single observation, weighted SSP and (adaptive) Bayes approaches for a travel distance of 12 m.

terrain transitions. In the weighted SPP approach, inerrance is caused by previous predictions which influence the final prediction significantly. Even if the current prediction is correct, misclassifications can occur if the posterior probability of the current prediction is low. The reduced reactivity in the Bayes approach can be explained using the chosen model parameter v . For the 0 and 0.6 m experiments the parameter value $v = 0.4$ is not optimal since, in some cases, the system erroneously decides to remain in the current state although a terrain transition has occurred. Here, a smaller value for v is preferable in order to increase the probability of a system state change.

For travel distances 1.2 and 2.4 m, the true positive rate of the adaptive Bayes approach is increasing. This indicates an appropriate choice of the parameter v . However, the choice is not optimal since the true positive rate is larger for the Bayes approach with a constant parameter of $v = 0.4$.

In comparison with the Bayes approach and the weighted SPP approach the adaptive Bayes technique yields increased true positive rates for travel distances larger than 2.4 m. Fig. 2(a) suggests that the adaptive Bayes approach automatically chooses a large value for v , which results in an increased model inerrance. This desired inerrance is larger than the one caused by the other two approaches, enabling the adaptive model to filter out classification errors occurring in an enlarged time frame. Experiments reveal a general trend of choosing large values for v . Already at a

travel distance of 6 m, v took values greater than 0.86 in 84.4% of all cases.

Considering the results of the second parameter set (Fig. 2(a)), the Bayes approach yields larger prediction errors in comparison with the single observation approach up to a travel distance of 1.2 m. Only for travel distances larger than 1.2 m the classification model benefits from exploiting temporal coherence. This behavior meets our expectations since we chose a large parameter value for v (0.85). As the probability of a system state change decreases with increasing v , the model becomes less reactive. However, in situations of high-frequent terrain transitions (as they appear in the 0, 0.6, and 1.2 m experiments) this is an unwanted property which results in a decrease of the true positive rate. With increasing travel distance terrain transitions become less likely. Thus, remaining in the current state is preferable, filtering out selective misclassifications. Now, the chosen value for v is more adequate since it reduces the probability of a system state change.

Considering the weighted SPP approach, the true positive rate is decreasing first, reaching its minimum at a travel distance of 1.2 m (72.5%), and is finally increasing with increasing travel distance. However, the classification errors occur more frequently as compared to the (adaptive) Bayes approach in all test cases. One reason for the decline in the true positive rate is the reduced reactivity in situations where

terrain transitions occur. In the worst case, the previous four classification steps correctly predict the previous terrain type, yet the most recent observation belongs to another surface ground. Not until several additional observations are recorded, enough data is available to extenuate the influence of previous predictions. Note that this effect is reduced as the travel distance increases due to the advantages arising from the filtering of misclassifications.

To evaluate whether the differences in the true positive rate are statistically significant, we performed two-tailed paired t-tests. Thereby, we compared the true positive rates of the 20 reruns of a single experiment using a certain travel distance d and the adaptive Bayes approach to those of another experiment employing another approach but the same travel distance. For the comparison we used all techniques and model parameter settings presented in Sect. IV-B. We performed a t-test of the null hypothesis that the 20 true positive rates are independent random samples from normal distributions with equal means and equal but unknown variances, against the alternative that the means are not equal. Thereby, we used a significance level of 5%. It turned out, that all differences of the true positive rate are statistically significant except for the ones when comparing the adaptive Bayes approach to the Bayes approach with parameter $v = 0.85$ for travel distances 12 and 16 m. Since the adaptive Bayes approach prefers similar values for v in comparison with the Bayes approach during model evaluation, these results follow our expectations.

Fig. 2(b) depicts the true positive rate for each individual terrain class. It shows that the decrease in erroneous predictions of the (adaptive) Bayes approach and the weighted SSP approach is based on an improved classification for all classes. Thereby, classification performance rises with the distance a robot navigates over the same terrain type. Comparing the adaptive Bayes and the single observation approaches in the 12 m experiment, we achieve an improvement of approximately 6.7% (indoor PVC floor), 5.6% (asphalt), 8.6% (gravel), 3.2% (grass), 11.0% (paving), and 9.8% (clay).

V. CONCLUSION

In this paper, we experimentally investigated the use of temporal coherence in the context of vibration-based terrain classification. We applied a Bayes filter which allowed for an effective inclusion of the classification history. However, only an adaptive approach which automatically adjusts its parameters was reactive enough to detect both, high-frequent and low-frequent terrain class changes. In comparison with previous SVM-based terrain classification approaches, which were based on single observations only, our adaptive Bayes filter approach yields an improvement in prediction performance of 1.3%, 4.2%, 6.3%, and 6.7%. These results denote the increase of classification performance when averaged over the complete set of terrain classes at travel distances of 2.4, 4, 8, and 16 m. We also compared our approach to

another technique which considered several observations for the final classification. The results showed the benefit of our Bayes approach which became more distinct as the distance between varying terrain classes increased.

Future research will focus on alternative classifiers to be embedded into our Bayes filter approach. Since the performance of our vibration-based terrain classification framework does not only depend on correct terrain classifications but also on the classification confidence we will address the choice of an appropriate classifier which provides reliable estimates of the posterior probability.

REFERENCES

- [1] B. H. Wilcox, "Non-geometric hazard detection for a mars micro-rover," in *Proceedings of the AIAA/NASA Conference on Intelligent Robots in Field, Factory, Service and Space*, vol. 2, Houston, TX, 1994, pp. 675–684.
- [2] R. Manduchi, A. Castano, A. Talukder, and L. Matthies, "Obstacle detection and terrain classification for autonomous off-road navigation," *Autonomous Robots*, vol. 18, pp. 81–102, 2005.
- [3] A. Angelova, L. Matthies, D. M. Helmick, and P. Perona, "Fast terrain classification using variable-length representation for autonomous navigation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Minneapolis, MN, USA, 2007, pp. 1–8.
- [4] N. Vandapel, D. Huber, A. Kapuria, and M. Hebert, "Natural terrain classification using 3-d lidar data," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2004)*, New Orleans, LA, April 2004, pp. 5117–5122.
- [5] J.-F. Lalonde, N. Vandapel, D. F. Huber, and M. Hebert, "Natural terrain classification using three-dimensional lidar data for ground robot mobility," *Journal of Field Robotics*, vol. 23, no. 10, pp. 839–861, 2006.
- [6] K. Iagnemma and S. Dubowsky, "Terrain estimation for high-speed rough-terrain autonomous vehicle navigation," in *Proceedings of the SPIE Conference on Unmanned Ground Vehicle Technology IV*, Orlando, FL, USA, 2002.
- [7] E. Coyle, E. G. Collins Jr, E. DuPont, D. Ding, H. Wang, R. A. Cooper, and G. Grindle, "Vibration-based terrain classification for electric powered wheelchairs," in *Proceedings of the IASTED Conference on Telehealth and Assistive Technologies*, Baltimore, USA, 2008.
- [8] E. M. DuPont, C. A. Moore, E. G. Collins Jr., and E. Coyle, "Frequency response method for terrain classification in autonomous ground vehicles," *Autonomous Robots*, vol. 24, no. 4, pp. 337–347, 2008.
- [9] E. M. DuPont, C. Moore, and R. G. Roberts, "Terrain classification for mobile robots traveling at various speeds: An eigenspace manifold approach," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2008)*, 2008, pp. 3284–3289.
- [10] C. Weiss, H. Fröhlich, and A. Zell, "Vibration-based terrain classification using support vector machines," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2006)*, Beijing, China, October 9–15 2006, pp. 4429–4434.
- [11] C. Weiss, N. Fechner, M. Stark, and A. Zell, "Comparison of different approaches to vibration-based terrain classification," in *Proceedings of the 3rd European Conference on Mobile Robots (ECMR 2007)*, Freiburg, Germany, 2007, pp. 7–12.
- [12] A. Yilmaz, O. Javed, and M. Shah, "Object tracking: A survey," *ACM Computing Surveys*, vol. 38, no. 4, pp. 13.1–13.45, 2006.
- [13] D. Fox, S. Thrun, F. Dellaert, and W. Burgard, "Particle filters for mobile robot localization," in *Sequential Monte Carlo Methods in Practice*. New York: Springer Verlag, 2000.
- [14] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, ser. Intelligent Robotics and Autonomous Agents. MIT Press, 2005.
- [15] C.-C. Chang and C.-J. Lin, *LIBSVM: a Library for Support Vector Machines*, Department of Computer Science and Information Engineering, National Taiwan University, Taipei 106, Taiwan, 2005.
- [16] T.-F. Wu, C.-J. Lin, and R. C. Weng, "Probability estimates for multi-class classification by pairwise coupling," *Journal of Machine Learning Research*, vol. 5, pp. 975–1005, 2004.
- [17] D. W. Scott and S. R. Sain, *Multi-Dimensional Density Estimation*. Amsterdam: Elsevier, 2004, pp. 229–263.