# Deep Spatiotemporal Models for Robust Proprioceptive Terrain Classification

•

# Abhinav Valada and Wolfram Burgard

#### Abstract

Terrain classification is a critical component of any autonomous mobile robot system operating in unknown realworld environments. Over the years, several proprioceptive terrain classification techniques have been introduced to increase robustness or act as a fallback for traditional vision based approaches. However, they lack widespread adaptation due to various factors that include inadequate accuracy, robustness and slow run-times. In this paper, we use vehicle-terrain interaction sounds as a proprioceptive modality and propose a deep Long-Short Term Memory (LSTM) based recurrent model that captures both the spatial and temporal dynamics of such a problem, thereby overcoming these past limitations. Our model consists of a new Convolution Neural Network (CNN) architecture that learns deep spatial features, complemented with LSTM units that learn complex temporal dynamics. Experiments on two extensive datasets collected with different microphones on various indoor and outdoor terrains demonstrate state-of-the-art performance compared to existing techniques. We additionally evaluate the performance in adverse acoustic conditions with high ambient noise and propose a noise-aware training scheme that enables learning of more generalizable models that are essential for robust real-world deployments.

#### Keywords

Convolutional neural networks, Recurrent neural networks, Long-short term memory networks, Terrain classification, Acoustics, Temporal classification

# 1. Introduction

As the transition of robots from mere lab equipment to autonomous machines that tackle complex problems in unknown environments occurs, the perceptual challenges that they face increases exponentially. Mobile robots in particular should have the capability to distinguish the terrain that they traverse on and accordingly optimize their navigation strategy. Failing to adapt the trafficability can lead to disastrous situations. For instance, a robot that can travel with a certain maximum speed on carpet, cannot use the same speed to traverse on sand, if not it would lead to entrenchment. Moreover, the predominately adopted vision based approaches have a high probability of failure due to the similar visual appearance of these two terrains. Often, there are also situations where terrains are covered with leaves or water which can be very challenging for visionbased classifiers.

Over the years, this has motivated researchers to explore alternate modalities such as lidars (Suger et al., 2015; Thrun et al., 2006), vibrations induced on the vehicles body (Brooks and Iagnemma, 2005), vehicleterrain interaction sounds (Christe and Kottege, 2016; Libby and Stentz, 2012; Ozkul et al., 2013; Valada et al., 2015) and roughness estimation using accelerometers (Eriksson et al., 2008; Weiss et al., 2006). Each of these approaches have their own benefits and drawbacks: optical sensors perform remarkably well in the presence of good illumination but are drastically affected by changes in lighting conditions. Classification with RGB images is usually achieved using color and texture features extracted from the scene (Sung et al., 2010). The use of texture-based image descriptors such as local ternary patterns have also been explored. In such approaches, features extracted using

Department of Computer Science, University of Freiburg, Germany http://deepterrain.cs.uni-freiburg.de

#### Corresponding author:

Abhinav Valada, Albert-Ludwigs-Universität Freiburg, Technische Fakultät, Autonome Intelligente Systeme, Georges-Köhler-Allee 079, 79110 Freiburg im Breisgau, Germany. Email: valada@cs.uni-freiburg.de local ternary patterns are used on sequences of mutated images and classified using Recurrent Neural Network (RNN) configurations (Otte et al., 2015). Lidars on the other hand, have extensively been used for traversability analysis. However, they are not suitable for fine-grained terrain classification where two or more terrains may have the same degree of coarseness. Nonetheless, they have been shown to benefit from being able to learn from partially labelled data and using semi-supervised learning techniques (Santamaria-Navarro et al., 2015; Suger et al., 2015). Commonly, features used include statistics on remission values, roughness and slope. As there are benefits in using both images and lidar data, approaches have also been explored to learn classifiers from both modalities. Here, features from images include color, texture, and geometric features and from lidar data typically surface normals, curvature, ground height, point feature histograms, linearity and planarity are used (Namin et al., 2014; Posner et al., 2008). CNN-based approaches for both near-range and farrange terrain classification using stereo and RGB images have been demonstrated during the Defense Advanced Research Projects Agency (DARPA) Learning Applied to Ground Robots (LAGR) program (Hadsell et al., 2009; Muller et al., 2013). These approaches achieved state of the art performance by combining both supervised and unsupervised learning using deep hierarchical networks. In contrast, acoustics based approaches perform well in all the above scenarios but are easily affected by environmental noise. Research using acoustics and other proprioceptive techniques are discussed in Section 2.

The benefit of exploring these alternate modalities is that the disturbances that affect optical or active sensors do not affect proprioceptive sensors such as microphones, hence enabling us to use them in a complementary classifier or fuse them with other modalities to increase robustness. These alternate approaches have failed to gain popularity beyond research. This can be attributed to the following: (1) manually designing feature sets that perform well in every real-world condition is tedious and impractical; (2) existing techniques have slow run times making them unusable for real-time applications; (3) most techniques require specific hardware setups that are difficult to replicate; and (4) approaches often lack reliability in real-world scenarios.

In this paper, we propose a novel multiclass terrain classification approach that overcomes these impediments by using only the sound from vehicle-terrain interaction to classify both indoor and outdoor terrains. We gathered two long-scale vehicle-terrain interaction datasets: one by equipping our mobile robot with a shotgun microphone and the other using a mobile phone microphone, as shown in Figure 1. The rationale behind this was to test the generalizability of the model to different



**Figure 1.** The Pioneer P3-DX platform that we use in our experiments, showing the shotgun microphone with the shock mount mounted close to the wheel.

hardware setups. Sounds from such interactions have distinct audio signatures and can be even used for fine-grained terrain classification such as distinguishing between grass and mowed-grass (Valada et al., 2015). Unlike speech or natural domestic sounds, vehicle-terrain interaction sounds are very unstructured in nature as several dynamic environmental factors contribute to the signal. Our previous work demonstrated that features learned using Deep Convolutional Neural Networks (DCNNs) significantly outperforms classification using shallow, manually-engineered feature sets. Our network combines three temporal pooling methods to achieve the time-series representation learning.

In this work, we extend our previous approach by incorporating a deep LSTM recurrent framework in order to further exploit the temporal dynamics in our problem. Recurrent neural models form compositional representations in the time domain, similar to convolutional models forming compositional representations in the spatial domain. LSTM based recurrent neural network models have recently achieved impressive results for sequential learning tasks such as language translation (Sutskever et al., 2014), speech recognition (Graves and Jaitly, 2014) and image captioning (Donahue et al., 2015). The structure of our proposed network leverages spatio-temporal information in each clip and in transitioning clips as well, where we define a clip as the time window that we take for classification. We show that our end-to-end trained recurrent neural model learns improved feature representations compared to our DCNN on both our datasets.

DCNNs learn highly discriminative features from training data, however, a caveat of this attribute is that the training data cannot encompass every possible realworld scenario. For acoustics-based terrain classification, this primarily relates to the ability of the model to adapt to ambient noise in the environment. Noise filtering techniques can be applied if the training data can include every possible noise disturbance, but this is highly impractical. Therefore, the model has to learn the distribution of common noises so that it can generalize effectively to different situations. Noise sensitivity is strongly correlated to the capability to operate in different real-world environments. Moreover, there are two different types of noise sources in our application; ambient noise and the noise from the motors of the robot. As the microphone is mounted close to the tires of the robot, it is inevitable that the captured signals also include the noise from the motors of the robot. Therefore this is already accounted for in the dataset. In order to make our model generalize effectively to real-world environments with varying types and amounts of ambient noise, we present a noise-aware training scheme. Our training scheme randomly injects ambient noise signals from the Diverse Environments Multichannel Acoustic Noise Database (DEMAND) (Thiemann et al., 2013) at different Signal-to-Noise (SNR) ratios. We then demonstrate that our recurrent model trained using our noise-aware scheme shows improved robustness in realworld conditions.

Specifically, the following are the main contributions of this paper:

- We propose a new deep spatio-temporal architecture for learning complex dynamics in proprioceptive signals.
- (2) We optimize various hyperparameters of the model and show their influence on the performance.
- (3) We demonstrate that our approach is hardware independent and the performance of our model on data from an inexpensive hardware is on par with data from a high-quality device.
- (4) We extensively evaluate the utility of our proposed Global Statistical Pooling (GSP) on various deep spatio-temporal model configurations.
- (5) We quantify our models performance in seven different environments having adverse acoustic conditions.
- (6) We introduce a noise aware training scheme that substantially increases the generalizability of the model to real-world scenarios.
- (7) We present thorough empirical evaluations on over 6h of audio data.

Rest of this paper is organized as follows. In Section 2, we discuss the work that has been previously done using acoustic and other proprioceptive sensors for terrain classification. We then detail our recurrent LSTM approach, the network architecture and training in Section 3. In Section 4, we provide a brief overview our data collection methodology. Finally we present our experimental evaluation in Section 5, followed by conclusions and discussion in Section 6.

#### 2. Related Work

Terrain classification using proprioceptive modalities has not been explored in the same depth as vision based approaches, yet there is a sizable amount of work in this area. The most researched proprioceptive terrain classification technique is using accelerometer data (Ojeda et al., 2006; Trautmann and Ray, 2011; Weiss et al., 2006). This is achieved by extracting features such as power spectral density, discrete fourier transform and other statistical measures from the vibrations induced on the vehicles body. Such approaches demonstrate a substantial amount of false positives for finer terrain classes such as asphalt and carpet. However, accuracies as high as 91.8% has been reported for a seven class problem using Support Vector Machines (SVMs) (Weiss et al., 2006). In a similar work, accelerometer data from a mobile sensor network system was used to detect potholes and other road anomalies (Eriksson et al., 2008). Hand engineered features were used and the system achieved an average accuracy of 90% in real-world experiments.

There is a body of work in terrain classification tailored to legged robots. Unlike in wheeled robots, proprioceptive terrain classification can enable safe foothold placement which is critical to ensure the stability of such legged systems. In one of the initial works (Hoepflinger et al., 2010), the authors extracted features from ground contact forces and joint motor current measurements to train a multiclass AdaBoost classifier. Although they did not test the efficacy of their approach on real-world terrains, they demonstrated the classifiers capability to identify different coarseness and curvatures of surfaces, as a first step towards real-world proprioceptive terrain classification. In another approach (Best et al., 2013), the authors demonstrate the ability to classify four different outdoor terrains using position measurements from leg servos of a hexapod robot. They extract a 600-dimensional feature vector consisting of gait-phase domain, as well as frequency domain features. Their approach utilizes a 2.7s window of data and a SVM is trained to classify the terrains.

Vibration data from contact microphones have also been successfully utilized for terrain classification. The vibrations captured are similar to accelerometers than air microphones that we use in this work, as they pick up only structure-borne sound and minimal environmental noise. Contact microphones have previously been mounted on analog rover's wheel frame to capture the induced vibrations and classify the terrain (Brooks and Iagnemma, 2005). They extract log-scale power spectral density features and train a pairwise classifier. For a three class problem, they achieve an average accuracy of 74% on a wheel-terrain testbed and 85.3% on a rover. They also demonstrate a self-supervised classification approach where a visual classifier is trained using labels provided by a vibration-based classifier (Brooks and Iagnemma, 2007).

The use of vehicle-terrain interaction sounds for terrain classification has been the most sparsely explored among all the proprioceptive modalities. Typically, combinations of state-of-the-art handcrafted audio features are used with traditional machine learning algorithms. Recently, a multiclass sound-based terrain classification system for a mobile robot was presented, that utilizes features including spectral coefficients, moments and various other spectral and temporal characteristics (Libby and Stentz, 2012). Their SVM based classifier achieves an average accuracy of 78% over three terrain classes and three hazardous vehicle-terrain interaction classes. They also showed that smoothing over larger temporal window of about 2s yields an improved accuracy of 92%. In another approach (Ojeda et al., 2006), a suite of sensors including microphones, accelerometers, motor current and voltage sensors, infrared, ultrasonics and encoders were used along with a feedforward neural network classifier. However their microphone based approach, only achieved an average accuracy of 60.3% for a five class problem. They concluded that the overall performance was poor and such an approach was promising only for classes such as grass.

Microphones have also been used for terrain classification with legged robots. In one such approach, terrain is classified from the steps taken by a hexapod robot (Ozkul et al., 2013). They use zero-crossing rate, frequency band coefficients and delta-features along with a functional trees classifier. They also experiment with different noise elimination techniques to remove motor/gear-head noise, but concluded that the performance was worse after the noise elimination. Their seven class classifier achieved an average accuracy of 90%. More recently, a SVM based system that uses statistics of spectral and band features, was used to classify terrain from the sounds produced during the locomotion of hexapod (Christe and Kottege, 2016). Their classifier was trained on a dataset that contains 5min worth of data from each terrain type. Their approach uses 1s windows for classification and operates at 1Hz. They report an average accuracy of 89.08% for seven terrain classes. They also investigated the use of spectral subtraction to eliminate the servo noise. They report an improvement in the sensitivity from 92.9% to 95.1% after spectral subtraction.

In all of the previous works, manually handcrafted features were extracted after specialized preprocessing steps. The approaches were evaluated on comparatively limited data and not in varying real-world environments. Most importantly, they do not model the temporally discriminative information in proprioceptive data which could potentially improve the classification performance, increase robustness and reduce false detections. In addition, they do not address the robustness of their models to ambient noise, which is one of the most critical properties for real-world adaptation. In the following section, we describe our DCNN architecture, the methodology using which we incorporate temporal recurrence into our model and our noise-aware training scheme.

#### 3. Technical Approach

One of the main objectives of our work is to develop an end-to-end trainable recurrent model tailored for classifying unstructured vehicle-terrain interaction sounds, with as minimal preprocessing as possible. Even though learning directly from raw waveforms has been demonstrated for speech processing, recent work has shown that the performance of such approaches is lower and the computational cost is higher, when compared to learning from simple transformations such as spectrograms (Graves and Jaitly, 2014; Khunarsal et al., 2013). Therefore we choose to train our models on top of this minimal preprocessing. This signal processing can also be computed in the input layer of the network. The features used for classification are learned by the network from this input representation. Figure 2 depicts our preprocessing pipeline.

# *3.1. Spectrogram Transformation and Augmentation*

Unlike speech decoding or handwriting recognition tasks that require specific alignment between the audio and transcription sequences, our application does not necessitate a specific target for each frame. We only need to classify a set of frames/clips without requiring any prior alignment in the sequences. Therefore, we use individual clips as a new sample for classification. We first split the raw audio signal into short clips of  $t_w$  and  $t_{ov}$  of overlap between the clips. We then experimentally identify the shortest clip length, overlap and the number of windows taken by the LSTM units that best enables our model to learn the discriminative information from the time varying vehicle-terrain interaction signal.

We extract the Short-Time Fourier Transform (STFT) based spectrogram for each clip in the sequence. We first block each audio clip into M samples with 50% overlap between the frames. Let x[n] be the recorded raw audio signal with duration of  $N_f$  samples,  $f_s$  be the sampling frequency, S(i, j) be the spectrogram representation of the 1-D audio signal and  $f(k) = kf_s/N_f$ . By applying STFT on the length M windowed frame of the signal, we get

$$X(i,j) = \sum_{p=0}^{N_f - 1} x[n] \ w[n-j] \exp\left(-p\frac{2\pi k}{N_f}n\right),$$
(1)  
$$p = 0, \dots, N_f - 1$$



invariant to certain signal perturbations.

We use a Hamming window function w[n] to compensate for the Gibbs effect while computing STFT by smoothing the discontinuities at the beginning and end of the audio signal, i.e.,

$$w[n] = 0.54 - 0.46 \cos\left(2\pi \frac{n}{M-1}\right),$$

$$n = 0, \dots, M-1$$
(2)

We then compute the log of the power spectrum as

$$S_{\log}(i,j) = 20\log_{10}(|X(i,j)|) \tag{3}$$

We chose  $N_f$  as 2048 samples, therefore the spectrogram contains 1024 Fourier coefficients. Experiments from our previous work (Valada et al., 2015) revealed that most of the spectral energy is concentrated below 512 coefficients, therefore we only use the lower 512 coefficients to reduce the computational complexity and runtime. The noise and intensity levels vary a fair amount in the entire dataset as we collected data in several real-world environments. Factors such as environmental noise at different times of the day, variations in weather conditions and variations in the servo noise at different speeds, contribute to this. Therefore, we normalize the spectrograms by dividing by the maximum amplitude. We then compute the mean spectrum over the entire dataset and subtract it from the normalized spectrogram. We compute this as

$$S(i,j) = S_{\log}(i,j) / \max_{i,j} S_{\log}(i,j)$$

$$\tag{4}$$

Data augmentation has been shown to improve the performance of several deep learning tasks and also help

in training of very deep networks. We experimented with several signal augmentation strategies and created additional samples by applying a set of augmentations  $A_t$  on the audio signal in the frequency domain. Using two dimensional affine transform and warping, random offsets in time and frequency were applied to perform shifting of the spectrogram. Furthermore, we created more samples using time stretching, modulating the tempo, random equalization, and by varying the volume gain. We also experimented with frequency and time normalization with a sliding window and local contrast normalization.

#### 3.2. Network Architecture

Recurrent Neural Network (RNN) models learn complex temporal dynamics from a sequence by mapping inputs  $x = (x_1, \ldots, X_T)$ , to a sequence of hidden states  $h = (h_1, \ldots, h_T)$ , and hidden states to an output sequence  $y = (y_1, \ldots, y_T)$ . This is achieved by iterating through the following equations from t = 1 to T, as

$$h_t = g(W_{ih}x_t + W_{hh}h_{t-1} + b_h)$$
(5)

$$y_t = W_{ho}h_t + b_o \tag{6}$$

where g is is the hidden layer activation function such as an element-wise application of sigmoid non-linearity, W and b terms denote the weight matrices and bias vectors, with subscripts i, h, and o denoting input, hidden and output respectively.

As discussed extensively in previous works (Donahue et al., 2015; Graves and Jaitly, 2014; Hochreiter and



Schmidhuber, 1997), traditional RNNs often suffer from vanishing and exploding gradient problems that occur from propagating gradients through several layers of the RNN. The larger the length of temporal input, the harder it is to train the RNN. The long short-term memory (LSTM) architecture (Hochreiter and Schmidhuber, 1997) was proposed as a solution to this problem and to enable exploitation of long-term temporal dynamics from a sequence. LSTMs incorporate memory units containing several gates that regulate the flow of information in and out of the cells.

We first use our DCNN architecture (Valada et al., 2015), to extract features from the spectrogram of each clip in the sequence in time. The spectrograms in our training set are of the form  $S = \{s_1, \ldots, s_T\}$  with  $s_i \in \mathbb{R}^N$ . Each of them are of size  $v \times w$  and number of channels d (d = 1 in our case). Our network, shown in Figure 3, consists of six convolution layers, six Cascaded Cross Channel Parametric Pooling (CCCP) layers (Lin et al., 2013), three Fully-Connected (FC) layers, an LSTM layer and softmax layer. All the convolution layers are one dimensional with a kernel size of three and convolve along the temporal dimension. We use a fixed convolutional stride of one. CCCP layers that follow the first, second and third convolution layers are used to enhance discriminability for local patches within receptive fields. CCCP layers effectively employ  $1 \times 1$  convolutions over the feature maps and the filters learned are a better nonlinear function approximator. A max-pooling layer with a kernel of two follows the second, fourth and sixth CCCP layers. Max-pooling adds some invariance by only taking the high activations from adjacent hidden units that share

the same weight, thereby providing invariance to small phase shifts in the signal.

DCNNs that are designed to operate on images for various perception tasks, specifically preserve the spatial information of features learned. However, our application does not benefit from this, as we are only interested in identifying the presence or absence of features, rather than localizing the features in the frame. To this end, we introduce a new Global Statistical Pooling (GSP) scheme that applies different pooling mechanisms across entire input feature maps and combines them to gather statistics of features across a particular dimension. Our architecture utilizes this GSP by incorporating three different global pooling layers after cccp-9 to compute the statistics of the features across time. We use an inner product layer to combine the outputs of max pooling, L2 norm pooling and average pooling. In our previous work (Valada et al., 2015), we investigated various combinations of different global pooling layers and found that for other configurations, the accuracy dropped over 3%. As our new proposed architecture utilizes LSTM units to model the temporal relationships, we explore the utility of this global statistics pooling in Section 5. We use the Xavier weight filler (Glorot and Bengio, 2010) to initialize the convolution and CCCP layers by drawing from a zero mean uniform distribution from [-a, a] and the variance as a function of input neurons  $n_{in}$ , where  $a = \sqrt{3 / n_{in}}$ . The effects of initializing with just a Gaussian filler are discussed in Section 5.

Rectified linear units (ReLUs) significantly help in overcoming the vanishing gradient problem. We use ReLUs  $f(x) = \max(0, x)$ , after the convolution layers and dropout regularization (Hinton et al., 2012) on the inner product



layers. We then stack the LSTM layer after fc12, followed by another inner product layer and a *softmax* layer. If  $x_t$  is the input to the LSTM layer at time t, the activations can be formulated as

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i) \tag{7}$$

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f) \tag{8}$$

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o) \tag{9}$$

$$g_t = \phi(W_{xc}x_t + W_{hc}h_{t-1} + b_c)$$
(10)

$$c_t = f_t \odot c_{t-1} + i_t \odot g_t \tag{11}$$

$$h_t = o_t \odot \phi(c_t) \tag{12}$$

where  $\sigma(x) = (1 + e^{-x})^{-1}$  is the sigmoid nonlinearity and  $\phi(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}} = 2\sigma(2x)$  is the hyperbolic tangent nonlinearity.  $h_t \in \mathbb{R}^N$  is the hidden unit,  $g_t \in \mathbb{R}^N$  is the input modulation gate,  $c_t \in \mathbb{R}^N$  is the memory cell,  $i_t \in \mathbb{R}^N$  is the input gate,  $f_t \in \mathbb{R}^N$  is the forget gate, and  $o_t \in \mathbb{R}^N$  is the output gate. W and b are the weight matrix and bias with subscripts i, f, h, c, and o representing input, forget, hidden, cell and output gates respectively.  $\odot$ represents element-wise multiplication. The hidden state  $h_t$ models the terrain that the robot is traversing on at time t. The output of the memory cell changes over time based on the past states and the current state of the cell. Therefore, the hidden state is formed based on the short-term memory of the past clip. At timestep t, the predicted distribution  $P(y_t)$ can be computed by taking the *softmax* over the outputs of the sequence from the LSTM units  $z_t$ , i.e.,

$$P(y_t = c) = \frac{\exp(W_{zc}z_{t,c} + b_c)}{\sum_{c' \in C} \exp(W_{zc}z_{t,c'} + b_c)}$$
(13)

We train the entire model end-to-end using minibatch Stochastic Gradient Decent (SGD) with a batch size of 256. We optimize the SGD by smoothing the gradient computation for minibatch using a momentum coefficient  $\alpha$  as  $0 < \alpha < 1$ . The update rule can then be written as

$$\Delta w_{ij}(t) = \alpha \Delta w_{ij}(t-1) - \epsilon \frac{\partial E}{\partial w_{ij}(t)}$$
(14)

We found the model to be extremely sensitive to the learning parameters used. Therefore, we use the Spearmint Bayesian optimization library (Snoek et al., 2012) to tune the hyperparameters. We first optimized the learning policy and the initial learning rate over fixed, inverse, step and poly policies. We obtained the best performance using an initial learning rate of  $\lambda_0 = 0.01$  and the poly policy as  $\lambda_n = \lambda_0 \times (1 - N/N_{\text{max}})^c$ , where N is the iteration number,  $N_{\text{max}}$  is the maximum number of iterations, and c is power. We tuned the number of outputs in the inner product layers and the LSTM layer as they were highly correlated and had a big impact on the performance of the model. We increased the number of parameters in the inner product layers in order to make the LSTM model converge. This parameter tuning is further discussed in Section 5.

# 3.3. Noise Aware Training

The basic principle of our noise aware training scheme is to randomly inject common environmental noises into the training data so that the signal features are learned along with noise patterns, and hence accounted for during inference. As our hierarchical model encompasses a spatial and a temporal component, the lower spatial layers learn features that describe both the pure signal and noise signal, whereas the higher layers learn to distinguish between them. The temporal component enables the model to learn the evolution of both the signals. Due to the spatio-temporal depth, the network is capable of learning heterogeneous signal patterns and at what stage to de-emphasize the noise while making decisions. Noise-conditioned decision boundaries are leaned as we train the network with signals corrupted with noise at different Signal-to-Noise Ratios (SNRs). The noise-aware training scheme also regularizes the network. It helps improve the classification of pure signals as the easily degraded parts of the signal are blurred by noise which forces the network to learn the more dominant features, hence avoiding over-fitting.

During the training stage, we randomly select noise following а multinomial distribution а  $Mult(\mu_1, \mu_2, \ldots, \mu_n)$ , where n is the types of noises and  $\mu_i$  is sampled from a Dirichlet distribution as  $(\mu_1, \mu_2, \ldots, \mu_n) \sim \text{Dir}(\alpha_1, \alpha_2, \ldots, \alpha_n)$ , where  $\alpha_i$  is set to control the base distribution of the noise types. Furthermore, the SNR of the noised sample follows a Gaussian distribution  $\mathcal{N}(\mu_{\text{SNR}}, \sigma_{\text{SNR}})$ , where  $\mu_{\text{SNR}}$  and  $\sigma_{\rm SNR}$  are the mean and variance. We randomly select a start point s on the noise signal and scale it according to a SNR before adding it to the pure vehicle-terrain clip. The noise samples are concatenated if the noise clip length is smaller than the vehicle-terrain clip length. We use the Praat framework (Boersma and Weenink, 2013) for the noise corruption.

We use 18 classes of noise recordings from the Diverse Environments Multichannel Acoustic Noise Database (DEMAND). We categorize the noise into seven classes for in-depth experiments as follows:

- White: White noise has a very wide band and it is one of the most common noise sources. It has a very similar effect to that of various physical and environmental disturbances including wind and water sources.
- Domestic: This category includes noises from living rooms, kitchens and washing rooms. As our terrain classes also contain indoor terrains, it is critical to train the model with common indoor ambient noises.
- Nature: The nature category contains outdoor noise samples from a sports field, a river creek with flowing water and a city park.
- Office: This category contains recordings from an office with people working on computers, a hallway where people pass by occasionally.and from interior public spaces.
- Public: The public category contains recordings from interior public spaces such a bus station, a cafeteria with people and at a university restaurant.
- Street: The street category contains noise recordings from outdoor inner-city public roads and pedestrian traffic. It contains a busy traffic intersection, a town square with tourists and a cafe at a public space.
- Transportation: This category contains recordings of vehicle noises such as cars, subways and buses.

As vehicle-terrain sounds are already unstructured in nature, corrupting it with various heterogeneous noises might lead to the model not converging. We first investigate the effect of these disturbances on models trained for audiobased terrain classification. We then train our network with a fixed SNR on all the above ambient noise categories and compare it with training on corrupted samples with varying SNRs. The weights and biases for the noise-aware training are initialized by copying them from our pure trained model as described in the previous section. We use a learning rate 1/10th of the initial rate used for training the network. Results from these experiments are discussed in section 5.4.

# 3.4. Baseline Feature Extraction

There is a wide spectrum of audio features developed for various audio recognition and detection tasks. In previous works, researchers have explored the utilization of these features for sound-based terrain classification (Christe and Kottege, 2016; Libby and Stentz, 2012; Ozkul et al., 2013). In order to evaluate the performance of our proposed model, we extracted several of these traditional baseline audio features. We explored the use of both time and frequency domain features. For wheeled mobile robots, previously, Ginna and Shape features have been demonstrated to yield the best results (Libby and Stentz, 2012). Ginna features (Giannakopoulos et al., 2006) is a six dimensional feature vector consisting of zero crossing rate (ZCR), short time energy (STE), spectral centroid (SC), spectral rolloff (SR) and spectral flux (SF). While, shape features (Wellman et al., 1997) is four dimensional feature vector consisting of spectral centroid, standard deviations, skewness and kurtosis.

ZCR can be defined as the rate at which the signal changes from positive to negative. ZCR is one of the most traditional audio feature used for speech and music classification. STE can be defined as the energy in short segment of the signal. It can be given as the square of the amplitude of the signal. Spectral centroid also referred to as the median of the power spectrum, indicates where the center of mass of the spectrum is. They can be calculated as

$$ZCR = \frac{1}{2N} \sum_{m=0}^{N-1} |\operatorname{sgn}[y(m)] - \operatorname{sgn}[y(i-1)]|$$

(15)

where, 
$$\operatorname{sgn}[x(m)] = \begin{cases} 1, x(n) \ge 0\\ -1, x(n) < 0 \end{cases}$$
 (16)

STE = 
$$\sum_{m=0}^{N-1} X((m).w(n-m))^2$$
 (17)

$$SC = \frac{\sum_{m=0}^{N-1} f(m)X(m)}{\sum_{m=0}^{N-1} X(m)}$$
(18)

where, X(n) is a sample of the power spectrum at bin n, f(n) is the center frequency of the bin n, w(n) is the windowing function, and N is the size of the window. Spectral flux indicates the change in the power spectrum between successive windows. Spectral rolloff is the frequency below the 95th percentile of the power in the spectrum. They can be given as

$$SF = \frac{\sqrt{\sum_{m=1}^{N-1} (X(m,n) - X(m,n-1))^2}}{N-1}$$
(19)

SR = K, where 
$$\sum_{m=0}^{K} X(m) = 0.95 \sum_{m=0}^{\text{fmax}} X(m)$$
 (20)



where fmax is the maximum frequency at bin m. Spectral skewness measures the symmetry of the distribution of the spectral magnitude around the mean. Spectral kurtosis measures the similarity of the distribution of the spectral magnitude to that of a Gaussian distribution. They can be defined by

SS = 
$$\frac{1}{S} \sum_{m=0}^{N-1} \left( \frac{X(m) - \mu}{\sigma} \right)^3$$
 (21)

$$SK = \frac{1}{S} \sum_{m=0}^{N-1} \left( \frac{X(m) - \mu}{\sigma} \right)^4 - 3$$
 (22)

where 
$$\mu = \frac{1}{S} \sum_{m=0}^{N-1} X(m)$$
 (23)

$$\sigma = \sqrt{\frac{1}{S} \sum_{m=0}^{N-1} (X(m) - \mu)^2}$$
(24)

where, S is half the size of the window S = N/2,  $\mu$  is the mean across the windows, and  $\sigma$  is the standard deviation across the windows. A combination of Melfrequency Cepstral Coefficients (MFCCs) and Chroma features have been demonstrated to yield state of the art performance for music classification tasks (Ellis, 2007). MFCCs are one of the most widely used features for audio classification and Chroma features are strongly related to the harmonic progression of the audio signal. We use a combination of 12-bin MFCCs and 12-bin Chroma features for our comparisons. Timbral features containing means and variances of spectral centroid, spectral rolloff, spectral flux, zero crossing rate, low energy, and first 5 MFCCs, have shown impressive results for music genre classification (Tzanetakis and Cook, 2002). We also use this 19 dimensional feature vector for our baseline comparison. For our final baseline comparison, we use a combination of cepstral features containing 13-bin MFCCs, Line Spectral Pair (LSP) and Linear Prediction Cepstral Coefficients (LPCCs) (Verma, 2008).

# 4. Data Collection and Labelling

To the best of our knowledge, audio-based terrain classification using mobile robots has been focused on outdoor terrains. Whereas, there has been some work on both indoor and outdoor audio-based terrain classification using legged robots. Our objective is to create a model that can classify a wide variety of both indoor and outdoor terrains. Some indoor and outdoor terrains can have very similar visual features (Figure 5(a), 5(h), 5(i)) and hence

Table 1. Classification accuracy of several baseline leature extraction approaches on our dataset for a clip length of 300ms. Our
recurrent neural model performs 8.83% better than the best baseline approach.

Features	SVM Linear	SVM RBF	k-NN
Ginna	$44.87\pm0.70$	$37.51\pm0.74$	$57.26 \pm 0.60$
Spectral	$84.48 \pm 0.36$	$78.65 \pm 0.45$	$76.02\pm0.43$
Ginna & Shape	$85.50\pm0.34$	$80.37\pm0.55$	$78.17\pm0.37$
MFCC & Chroma	$88.95 \pm 0.21$	$88.55 \pm 0.20$	$88.43 \pm 0.15$
Trimbral	$89.07 \pm 0.12$	$86.74 \pm 0.25$	$84.82\pm0.54$
Cepstral	$89.93 \pm 0.21$	$78.93 \pm 0.62$	$\textbf{88.63} \pm \textbf{0.06}$
DCNN (ours)		$97.52 \pm 0.016$	
DCNN with LSTM (ours)		$\textbf{98.76} \pm \textbf{0.009}$	

pose a challenge for classifying using the vision based counterparts. We collected over 15h of vehicle-terrain interaction data from six different outdoor terrains and three different indoor terrains.

Our approach is platform independent, as the network learns features from the given training data. For our experiments, we use the P3-DX mobile robot platform to collect data as it has a small footprint and feeble motor noise. We equipped the P3-DX with rugged wheels that enabled us to collect data both indoors and structured as well as unstructured outdoors environments. One of the big challenges for audio classification is to deal with ambient and environmental noise. Training the network with data that includes interference from nearby sound sources can drastically affect the performance. In order to prevent such biases in training data, we equip the robot with a shotgun microphone to record vehicle-terrain interaction sounds. A shotgun microphone has a supercardioid polar pickup pattern, which helps in rejecting off-axis ambient sounds. Specifically, we chose the Rhode VideoMic Pro and mounted it near the wheel of the mobile robot as shown in Figure 1. The microphone has an integrated shock mount that prevents the pickup of any unwanted vibrations caused during the traversal.

We collected data in several different locations to have enough variability and to enable our model to generalize effectively to real-world environments. Therefore even signals in the same class have varying spectral and temporal characteristics. The robots speed was varied from  $0.1 \mathrm{m \, s^{-1}}$ to  $1.0 \mathrm{m \, s^{-1}}$  during the data collection runs. The data was recorded in the lossless 16-bit WAV format at 44.1kHz to avoid any recording artifacts. There was no software level boost added during the recordings as we found it to amplify the ambient noise substantially, instead we used a 20dB hardware level boost. The data was manually labelled using the live tags and timestamps made during the recordings. A waveform analyzer tool was used to crop out significant noise disturbances that had substantially higher amplitudes than the vehicle-terrain interaction signals. By

noisy disturbances, we refer to uncommon temporary environmental disturbances such as a nearby car or train passing. The data is then split to train and test sets, ensuring that the classes have approximately the same number of samples to prevent any bias towards a specific class. It was also ensured that the training and validation sets do not contain clips from the same location.

# 5. Experimental Results

We use the Caffe (Jia et al., 2014) deep learning framework for our implementations and the LSTM described in (Zaremba and Sutskever, 2014) but a version of the implementation faster than (Donahue et al., 2015). All our models are trained end-to-end and the experiments were run on a system containing a NVIDIA TITAN X GPU with cuDNN acceleration. The results from our experiments are described in the following sections.

# 5.1. Baseline Comparisons

In this section, we empirically evaluate our base DCNN model and our LSTM model with several baseline audio features discussed in Section 4.3. For all the baseline experiments that we present in Table 1, we choose a fixed clip length of 300ms. We compare with SVMs and k-Nearest Neighbours (kNNs) classifiers, which have shown the best performance in the work by (Libby and Stentz, 2012). SVMs perform well in high dimensional spaces and kNNs perform well when there are irregular decision boundaries. As a preprocessing step for the baseline features, we normalize the data to have zero mean. We use the one-vs-all voting scheme with SVMs to handle multiple classes and experimented with Linear and Radial Basis Function (RBF) as decision functions. We use inverse distance weighting for kNNs and optimized the hyperparameters for both classifiers by a grid-search using cross-validation. For the baseline classifiers, we used the implementations in scikit-learn and LibSVM. The results from this comparison are showed in Table 1.



The best performing baseline feature combination was Cepstral features using a linear SVM kernel, achieving an accuracy of 89.93%. The performance of Trimbral features using a linear SVM kernel is comparable. Ginna and Shape features using an SVM RBF kernel was the best performing feature set in the work of (Libby and Stentz, 2012). It achieved an accuracy of 80.37%, which is 9.56% less than our best performing baseline. The worst performance was using only Ginna features with an SVM RBF kernel. Table 1 shows that the features containing MFCCs outperform other feature combinations. The feature sets used also contain the mean and standard deviations of each feature.

Our DCNN yields an accuracy of 97.52%, which is an improvement of 7.59% over the best performing Cepstral features and 12.02% over Ginna and Shape features. This is currently the state of the art for acoustics-based terrain classification for a clip length of 300ms. Our recurrent LSTM model achieves an improved accuracy of 98.76% with the same clip length and an LSTM window of 3. This demonstrates that learning spatio-temporal relationships can further improve the classification performance. Note that this is not the best result of our network; model analysis and parameter tuning, which further improve the model, are discussed in the following sections. Our DCNN model performs inference in 9.15ms and our recurrent LSTM model performs inference in 12.37ms for a clip length of 300ms, whereas baseline approaches presented in Table 1 have feature extraction and classification time in the order of a few seconds.

Moreover, Figure 6 shows the comparison of error rates of our DCNN model and DCNN with LSTM model. For all the classes other than Cobble, there is a considerable decrease in the error rate. For the Cobble class, there is an insignificant increase of 0.0008% in error. Paving and Offroad classes have the largest decrease in error. In our previous work, we showed that these two classes have the highest false positives and the results now demonstrate that even this complex relationship can be learned using our temporally recurrent model.

#### 5.2. Network Parameter Estimation

In this section, we describe our experiments that we perform to gain insight on the effect of learning spatiotemporal relationships using different DCNN model configurations. We first describe each of the models and then present an in-depth analysis of their performance. For consistency, we use the same 300ms clip length, similar to the baseline experiments.

We consider the following four spatio-temporal model configurations:

- (1) M1 (DCNN): This model has a structure similar to standard classification models where there are alternating convolution and pooling layers which help in learning spatial features across the frames of a clip, followed by an inner product and softmax layer. The convolutions are only one dimensional and convolve along the temporal dimension.
- (2) M2 (DCNN with GSP): This model is similar to the M1 model described above but it incorporates our global statistical pooling before the inner product layer, which helps in learning statistics of temporal features across time.
- (3) M3 (DCNN with LSTM): This model is a temporal extension of the M1 model described above. An LSTM layer is appended after the first two inner product layers to learn complex temporal dynamics across clips. The LSTM layer is followed by one inner product layer and a softmax layer.
- (4) M4 (DCNN with GSP & LSTM): This model is a temporal extension of the M2 model described above. Global statistical pooling is first used to combine statistics of pooled features across time, then the resulting features are fed to an LSTM module to learn temporal dynamics across several clips. Similar to M3, the LSTM layer is followed by one inner product layer and a softmax layer. We call this model DCNN ST (Spatio-temporal) for future comparisons.

We use the same train and test splits for the 300ms clips as in Valada *et al.*, 2015. For the LSTM models, we use

**Table 2.** Comparison of our deep spatio-temporal model configurations. Our temporally recurrent model demonstrates 7.47% improvement compared to its spatial DCNN counterpart.

Model	Acc.	Prec.	Rec.
M1 (DCNN)	91.29	91.88	91.56
M2 (DCNN + GSP)	97.52	97.56	97.61
M3 (DCNN + LSTM)	95.73	95.93	95.88
M4 (DCNN + GSP & LSTM)	<b>98.76</b>	98.75	98.82

 
 Table 3. Classification accuracy of our DCNN ST model for varying LSTM window lengths and audio clip lengths.

 Accuracy in %.

Window Size	2	3	4	5	6
300ms	97.43	98.76	80.09	98.70	98.62
$250 \mathrm{ms}$	96.74	97.38	98.93	79.96	98.67
$200 \mathrm{ms}$	96.69	96.22	97.86	99.03	79.58

a window size of three. Results from this comparison are shown in Table 2. Interestingly, the DCNN with GSP model outperforms the DCNN with LSTM model by 1.79%. This can be attributed to the fact that the LSTM model does not benefit much from learning using large clip lengths and short temporal windows. In the later part of this section, we investigate the influence of clip lengths and temporal window lengths. However, the DCNN with GSP and LSTM outperforms all the other models by achieving an accuracy of 98.76%, which is a 7.47% improvement over the performance of the DCNN model. This also illustrates that GSP is critical for learning effective spatio-temporal relations in vehicle-terrain interaction sounds.

Figure 7 shows the per-class sensitivity of the spatiotemporal models discussed. The performance of the DCNN with GSP and LSTM model surpassed all the other models for every class other than Grass Medium-High. The performance of the DCNN with GSP model is better than that of the DCNN with GSP and LSTM model for the Grass Medium-High class. All the models have near perfect sensitivity for the Carpet class. This is primarily because the spectral responses for this class is mostly flat while compared to the others. The LSTM models perform better for flatter terrains such as Wood, Carpet and Paving, while the GSP models perform better when the terrains are more irregular. The DCNN model without the temporal features perform substantially worse than the others for almost all the classes.

Moreover, there are a number of critical parameters that are required to be optimized for efficient spatio-temporal learning. The biggest trade-off is with selecting the optimal clip length. As each clip is a new sample for classification,



spatio-temporal model configurations. Our temporally recurrent model outperforms the other configurations in almost all the classes.

Table 4. Performance comparison of our DCNN model (M2) and DCNN ST model (M4) at varying audio clip lengths and the corresponding time taken to process though the pipeline. Accuracy in % and time in ms.

Clip Length	DCN	N	DCNN ST					
	Accuracy	Time	Accuracy	Time				
2000 ms	99.86	45.40	99.88	23.27				
1500 ms	99.82	34.10	99.83	21.16				
1000 ms	99.76	21.40	99.78	16.43				
500 ms	99.41	13.30	99.45	13.75				
300 ms	97.36	9.15	98.76	12.37				
250 ms	94.05	9.15	98.93	12.36				
200 ms	91.30	9.14	99.03	12.23				

the shorter the clip length, the higher is the rate at which we can infer the terrain. A larger clip length yields increased accuracy, although this leads to an increase in the execution time and decrease in the classification rate, which are both undesirable. For our application, fast classification and execution rates are essential for making quick trafficability decisions. In our previous work (Valada et al., 2015), we showed that a clip length of 300ms was sufficient considering the above trade-offs. However, the trade-off now has another degree of complexity, as the LSTM window size also needs to be optimized. Nonetheless, using LSTMs to learn temporal relationships allows us to use a shorter clip length. We investigate the relationship between these parameters by individually training models with various clip lengths and LSTM window lengths. Table 3 shows the results from these experiments.



The results show that the accuracy does not increase just by considering a larger LSTM window length. In fact, we found the models to be increasingly difficult to train for large window lengths and it was difficult to ascertain when they actually converge even if we train them for twice the number of epochs as our base model. For a 300ms clip, we obtained the best performance using a LSTM window size of three, above which the accuracy dropped significantly for a window size of four, and saturated for higher window lengths. As we experimented with lower clip sizes, we found the same pattern of obtaining the best performance for a certain window length and then a significant drop, followed by saturation for higher lengths. We also trained models with lower clip length than 200ms and with higher LSTM window lengths but they did not yield a better performance than the reported models in Table 3. The overall best performance for various LSTM window lengths and clip sizes was obtained using a clip length of 200ms and with a window size of five.

Furthermore, we compared the performance of our DCNN model (M2) and our DCNN ST model (M4) at varying audio clip lengths. Results from this comparison are shown in Table 4. It can be seen that our DCNN ST model outperforms our base M2 model for all audio clip lengths between 200ms and 2000ms. Furthermore, the DCNN ST model also has faster classification rates for almost all clip lengths. The speed-up is achieved by computing gradients with respect to recurrent weights in a single matrix computation. Figure 8 shows the classification rates of both the models at varying audio clip lengths. Although the DCNN model has a faster classification rate

for a clip length of 200ms, the accuracy is 7.73% lower than that of the DCNN ST model. Considering the clip length to classification time trade-off, we choose the 200ms DCNN ST model for the in-depth performance evaluation experiments described in the following sections.

# 5.3. Performance Evaluation

To further investigate the performance, we computed the confusion matrix, which gives us insights on the misclassifications between the classes. Figure 9(b) shows the confusion matrix of our best performing DCNN ST model for a clip length of 200ms and an LSTM window size of five. This corresponds to a substantial improvement in the misclassifications compared to our previous DCNN model (M2). Figure 9(a) shows the confusion matrix of our M2 model. Paving and Offroad, which had the highest misclassification rate previously, now has a 3.15% decrease in the false positives. The second highest rate of false positives was between Wood and Linoleum. Our current DCNN ST model demonstrates a 1.66% decrease in the false positives between Wood and Linoleum. There is also a 0.88% decrease in the misclassification between Offroad and Grass Medium-High. The misclassifications are primarily due to very similar spectral responses between the classes for short clip lengths.

The best performing classes were Carpet and Linoleum. In fact, the Carpet class demonstrated no false positives. All the classes other than Offroad, Wood and Grass Medium-High yielded accuracies in the high 99% range. By incorporating LSTM units, our new proposed model outperforms our previous state-of-the-art model by 1.51% and with a much smaller clip length. The per-class recall of this model is shown in Figure 9(d). This gives us insights on the ratio of correctly classified instances. The class with the lowest recall was Wood, whereas, the class with the highest recall was Carpet. The overall recall of the DCNN ST model was 99.05%.

#### 5.4. Robustness to Noise

In this section, we detail our experiments on the noise corrupted samples from the DEMAND noise database categorized into seven classes and at varying SNRs as described in Section 3.3. We first only added noise to our test set using methods described in (Loizou, 2007) and experimented on the model trained on clean noise-free signals. Table 5 shows detailed results from this experiment. We observe that the performance drastically decreases with decreasing SNRs. Moreover, different ambient noises affect the performance at varying degrees. Figure 10 shows the accuracy of our model for the various noises. It can be seen that for all the classes the models are fairly robust at SNRs until 20dB, thereafter the performance drops rapidly for some classes such as noises from the Domestic and



(a) Confusion matrix of our DCNN model (M2) for a clip length of 300ms (Valada et al., 2015)



(c) Per-class recall of our DCNN model (M2) for a clip length of  $300\mathrm{ms}$  (Valada et al., 2015)



(b) Confusion matrix of our DCNN ST model (M4) for a window length of five and clip length of  $200 \mathrm{ms}$ 



(d) Per-class recall of our DCNN ST model (M4) for a window length of five and clip length of  $200\mathrm{ms}$ 



Street categories, as well as White noise. This can be attributed to the broadband of these noises, which corrupts signals in most frequencies, while the other types of noises are present only at specific frequencies. Hence, vehicleterrain interaction signals consisting of high frequency components are relatively robust to noises such as in the category of Nature due to their presence mostly only in the low frequencies.

Ambient noises in the domestic category have the most damaging effect, where the mean accuracy at SNRs from 30 dB to -10 dB was 61.71%. It can be seen in Figure 10 that the model has a linear drop in accuracy for

decreasing SNRs, for noises from the Domestic category. The second most damaging noises were from the Street category, achieving a mean accuracy of 68.36%. The model demonstrates remarkable robustness to ambient noises from the Nature category, for which the model yielded an average accuracy of 97.49%, even for such low SNRs. This is primarily because noises in the Nature category have very distinct spectral patterns that have dominant structures, unlike vehicle-terrain interaction signals, hence the model is inherently robust to these corruption patterns.

It is also of interest to know how each of these noises influence the classification of specific terrains. Therefore

Noise	SNR (dB)													
	Clean	30	20	10	0	-10	mean							
White	99.03	99.00	93.42	69.66	20.85	9.16	58.42							
Domestic	99.03	99.03	99.03	99.03	99.03	99.03	99.03	99.03	98.63	82.84	55.24	25.38	9.16	54.25
Nature	99.03	99.03	99.02	98.99	98.63	90.27 22.23 12.99	97.19							
Office	99.03	99.03	99.02	98.65	77.40		79.27							
Public	99.03	99.02	99.01	98.17	73.01		76.44							
Street	99.03	98.93	95.23	71.10	36.19	9.68	62.22							
Transportation	99.03	99.03	98.99	97.48	57.02	28.65	76.23							
mean	99.03	98.95	95.36	84.18	60.87	26.02								

**Table 5.** Influence of various ambient noises on the classification accuracy of our DCNN ST model (M4). Noise samples are extracted from the DEMAND noise database. Noises in the White, Domestic and Street categories are the most damaging, whereas, noises in the Nature category are the least damaging.

**Table 6.** Influence of various ambient noises on the classification accuracy of our noise-aware DCNN ST model. There is a substantial increase in the robustness in all the noise categories.

Noise	SNR (dB)												
	Clean	30	20	10	0	-10	mean						
White	99.72	99.68	98.77	97.63	97.11	96.38	97.91						
Domestic	99.72	99.66	98.69	97.87	97.04	96.97	98.05						
Nature	99.72	99.71	99.69	99.51	99.08	97.28	99.05						
Office	99.72	99.72	99.65	98.70	98.61	96.00	98.53						
Public	99.72	99.70	99.67	98.71	98.49	97.63	98.84						
Street	99.72	99.68	98.73	97.90	97.81	96.36	98.1						
Transportation	99.72	99.69	98.77	98.62	97.65	97.47	98.44						
mean	99.72	99.69	99.13	98.42	97.97	96.87							



we computed the per-class precision for each of the noise categories. Results from this experiment are shown in Figure 11 for ambient noises from the DEMAND noise database and in Figure 12 for White noise. The Carpet class is recognizable for all the ambient noises and at all SNRs. The Paving class is the most affected by noise corruption, followed by Offroad and Asphalt. Interestingly, classification of indoor terrains such as Linoleum, Wood and Carpet are less affected by indoor noises than outdoor terrains. Although the converse does not appear to be true, where the outdoor terrains are less affected by outdoor noises. This is due to the reason that, generally, outdoor ambient noises have continuous noise corruption for long periods, whereas indoor ambient noises are usually short impacts. A curious phenomenon that can be observed in certain terrain classes for noise corruption with SNRs of -10dB with ambient noises such as Office, Public, Street and Transportation, is the sudden increase in precision compared to higher SNRs. This is primarily due to the fact that the spectral responses of these noises blur the inherently noisy lower frequency components of the vehicle-terrain interaction signals which enhances the ability of the network to classify these terrains by focusing on the more distinct higher frequency responses.



Figure 11. Per-class precision of our DCNN ST model when subject to various noises from the DEMAND noise database at different SNRs. Terrains are still recognizable for ambient noises in the Nature category even for SNRs lower than 0dB.



White noise on the other hand severely affects the classification ability below 0dB. Terrain classes such as Mowed-grass, Asphalt, Cobble and Wood are robust to noise corruption upto 0dB. Whereas, classes such as Linoleum, Offroad, Grass Medium-high and Carpet have an exponential increase in classification error for decreasing SNRs. Paving and Offroad terrains are not recognizable for SNRs greater than 10dB. These experiments demonstrate the need for noise robustness in order to operate in high ambient noise environments.

As described in Section 3.3, we performed noise-aware training by injecting ambient noises from the DEMAND noise database at varying SNRs. The noise injection helps the network learn probable noise corruption patterns that may occur in real-world environments. Results from this experiment are shown in Table 6. It can be observed that there is a substantial increase in classification accuracy for SNRs lower than 10dB. Even for very low SNRs, the model demonstrates state-of-the-art performance. The noise-aware model shows an improved performance even for pure signals without noise corruption. A comparison of the error rates of our base DCNN ST model and our noise-aware DCNN ST model for each of the noise categories in the DEMAND noise database is shown in Figure 12. Our base model demonstrated varying patterns of error change for increasing SNRs, whereas our noiseaware model shows a very gradual exponential increase for all the noise categories. This shows that noise-aware training enables the network to learn the general distribution of noise corruption for all the SNRs and for varying types of noises as well.

For the second part of the robustness experiments, we tested our model on our mobile phone microphone dataset. This dataset was collected at a different location than our main dataset on which we trained our models. Each of the samples in this dataset was tagged with a GPS location for visualizing the trajectory in maps and to correlate to the terrain. This dataset contains about 2.5h of audio. Unlike the shotgun microphone that we used for collecting our main dataset, mobile phones have a condenser microphone that collects sound from every direction, thereby also collecting considerable amount of ambient noise. The main purpose of testing on this dataset was threefold: (1) to verify the adaptability of the model to a new environment; (2) to quantify the models performance with a new hardware; and (3) to quantify the performance in the presence of substantial amount of real-world ambient noise.

Figure 14 shows an example trajectory that the robot traversed during experimentation. The figure also shows the variation in speed (0-2ms) along the path. The speed at which the robot traversed with along the path is indicated with red lines. Thicker the red lines, slower is the speed. Our model achieves an accuracy of 98.62%on this dataset. Figure 15 shows the confusion matrices of our DCNN model and DCNN ST model tested on the mobile phone microphone dataset. The classes that show the largest misclasifications using our DCNN ST model are Paving, Wood and Linoleum. There is about 3% decrease in misclassifications between Paving and Offroad in our DCNN ST model, while compared to the DCNN model. There is also an equivalent decrease in misclassification between the Cobble and Offroad classes. Compared to our DCNN model, the DCNN ST model shows an increase in misclassification between Wood and Linoleum, however, the overall accuracy of the DCNN ST model is still higher than our previous DCNN model. These experiments have demonstrated the utility of noise-aware training, making it a necessary step for robust terrain classification in noisy realworld environments.

#### 6. Conclusion

In this work, we introduced a robust proprioceptive terrain classification system based on recurrent convolutional neural networks, which uses only sound from vehicleterrain interactions to classify a wide range of indoor and outdoor terrains. The performance of our models surpass several baseline approaches, achieving state-ofthe-art results in proprioceptive terrain classification. Our proposed model is both spatially and temporally deep, and our results demonstrate that learning temporal dynamics can improve classification than when learning only in the spatial domain. We investigated the influence of various hyperparameters on the performance of our network and tuned them to obtain the best classification rate to accuracy



ratio. Our implementation operates on a 200 ms window at 81.7 Hz and achieves an overall accuracy of 99.03%.

Additionally, we thoroughly evaluated the robustness of our models to various extreme ambient noises and introduced a noise-aware training scheme that injects random noise into the training data that increases the overall robustness to a great extent. Our noise-aware model achieves an improved accuracy of 99.72% and the extensive experiments demonstrate the capability of our model to adapt to real-world environments with different ambient noises. We also presented empirical evaluations with an inexpensive low-quality microphone that shows the hardware independence and robustness of our approach.

#### Funding

This work has been partly supported by the European Commission under the grant numbers ERC-AGPE7-267686-LifeNav and FP7-610603-EUROPA2.

#### References

- G. Best, P. Moghadam, N. Kottege, and L. Kleeman. Terrain classification using a hexapod robot. In Australasian Conference on Robotics and Automation, 2013.
- P. Boersma and D. Weenink. Praat: doing phonetics by computer [computer program]. In Version 5.3.51, retrieved 2 June 2013 from http://www.praat.org/, 2013.
- C. A. Brooks and K. Iagnemma. Vibration-based terrain classification for planetary exploration rovers. *IEEE*

Mow Grass	-99.44	0.15	0.00	0.03	0.00	0.16	0.01	0.20	0.00	100 90	4ow Grass -	99.13	0.04	0.04	0.08	0.08	0.46	0.04	0.13	0.0		100 90
Paving	0.03	92.73	0.00	0.01	0.25	6.60	0.03	0.35	0.00 -	- 80	Paving -	0.43	95.49	0.0	0.06	0.31	3.69	0.01	0.02	0.0	-	80
Linoleum	0.00	0.00	99.95	0.00	0.00	0.00	0.05	0.00	0.00	- 70	Linoleum -	0.01	0.0	99.75	0.0	0.0	0.0	0.24	0.0	0.0		70
Asphalt	0.00	0.04	0.00	99.90	0.01	0.00	0.06	0.00	0.00	- 60	Asphalt -	0.02	0.08	0.01	99.81	0.01	0.01	0.06	0.0	0.0	-	60
Cobble	0.00	0.04	0.00	0.01	97.36	0.30	0.00	2.26	0.02	- 50	Cobble -	0.02	0.16	0.0	0.02	99.1	0.38	0.0	0.31	0.0	-	50
Offroad	0.01	0.48	0.00	0.00	0.03	98.43	0.00	1.04	0.00	- 40	Offroad -	0.05	0.2	0.0	0.0	0.17	99.2	0.0	0.38	0.0	-	40
Wood	0.00	0.01	0.09	0.05	0.01	0.00	99.57	0.04	0.23	- 30	Wood -	0.04	0.02	2.5	0.44	0.0	0.0	96.98	0.0	0.02	1	30
Grass MH	0.00	0.00	0.00	0.00	0.10	0.27	0.00	99.59	0.04	- 20	Grass MH -	0.09	0.01	0.0	0.0	0.43	0.72	0.0	98.65	0.09		20
Carpet	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00	- 10	Carpet -	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0		10
NOW CLOSE PAIRS CADE OFFOR WOOD STATE OF CONTRACT OF CADE OFFOR CA																						
<ul><li>(a) Confusion matrix of our DCNN model (M2) for a clip length of 300ms (Valada et al., 2015)</li><li>(b) Confusion matrix of our DCNN ST model (M4) for a window length of five and clip length of 200ms</li></ul>												/										

**Figure 14.** Comparison of our new DCNN ST model with our DCNN model on the mobile phone dataset. There is a significant decrease in misclassifications between Paving and Ofroad classes, as well as Cobble and Grass Medium-High.



**Figure 15.** Map showing one of the trajectories that the robot followed during the classification experiments using a mobile phone microphone. Variation in speed along the path is shown in red. Thicker red lines denote slower speed.

Transactions on Robotics, 21(6):1185–1191, Dec 2005.

- C. A. Brooks and K. D. Iagnemma. Self-supervised classification for planetary rover terrain sensing. In *Aerospace Conference*, 2007 IEEE, pages 1–9, March 2007.
- J. Christe and N. Kottege. Acoustics based terrain classification for legged robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2016.

- J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.
- D. P. W. Ellis. Classifying music audio with timbral and chroma features. In 8th International Conference on Music Information Retrieval, 2007.
- J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan. The pothole patrol: Using a mobile sensor network for road surface monitoring. In *The Sixth Annual International conference on Mobile Systems, Applications and Services (MobiSys 2008)*, Breckenridge, U.S.A., June 2008.
- T. Giannakopoulos, D. Kosmopoulos, A. Aristidou, and S. Theodoridis. Advances in Artificial Intelligence: 4th Helenic Conference on AI, SETN 2006, Heraklion, Crete, Greece, May 18-20, 2006. Proceedings, chapter Violence Content Classification Using Audio Features, pages 502–507. Springer Berlin Heidelberg, 2006.
- X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *In Proceedings of* the International Conference on Artificial Intelligence and Statistics (AISTATS10). Society for Artificial Intelligence and Statistics, 2010.
- A. Graves. Generating sequences with recurrent neural networks. arxiv preprint arxiv: 1308.0850, 2013.
- A. Graves and N. Jaitly. Towards end-to-end speech recognition with recurrent neural networks. In T. Jebara and E. P. Xing,

editors, Proceedings of the 31st International Conference on Machine Learning (ICML-14), pages 1764–1772, 2014.

- R. Hadsell, P. Sermanet, J. Ben, A. Erkan, M. Scoffier, K. Kavukcuoglu, U. Muller, and Y. LeCun. Learning longrange vision for autonomous off-road driving. *Journal of Field Robotics*, 26(2):120–144, Feb. 2009. ISSN 1556-4959.
- G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, Nov. 1997.
- M. A. Hoepflinger, C. D. Remy, M. Hutter, L. Spinello, and R. Siegwart. Haptic terrain classification for legged robots. In 2010 IEEE International Conference on Robotics and Automation, pages 2828–2833, May 2010.
- Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. arXiv preprint arXiv:1408.5093, 2014.
- P. Khunarsal, C. Lursinsap, and T. Raicharoen. Very short time environmental sound classification based on spectrogram pattern matching. *Information Sciences*, 243:57 – 74, 2013.
- J. Libby and A. T. Stentz. Using sound to classify vehicle-terrain interactions in outdoor environments. In *Proceedings of the IEEE International Conference on Robotics and Automation*, May 2012.
- M. Lin, Q. Chen, and S. Yan. Network in network. arXiv preprint arXiv:1312.4400, 2013.
- P. Loizou. Speech Enhancement: Theory and Practice. Signal processing and communications. 2007.
- U. A. Muller, L. D. Jackel, Y. LeCun, and B. Flepp. Real-time adaptive off-road vehicle navigation and terrain classification. *Proc. SPIE*, 8741:87410A–87410A–19, 2013.
- S. T. Namin, M. Najafi, and L. Petersson. Multi-view terrain classification using panoramic imagery and lidar. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 4936–4943, Sept 2014.
- L. Ojeda, J. Borenstein, G. Witus, and R. Karlsen. Terrain characterization and classification with a mobile robot. *Journal of Field Robotics*, 23(2):103–122, 2006.
- S. Otte, S. Laible, R. Hanten, M. Liwicki, and A. Zell. Robust visual terrain classification with recurrent neural networks. In *In Proc. of European Symposium on Artificial Neural Networks*, 2015.
- M. C. Ozkul, A. Saranli, and Y. Yazicioglu. Acoustic surface perception from naturally occurring step sounds of a dexterous hexapod robot. *Mechanical Systems and Signal Processing*, 40(1):178–193, 2013.
- I. Posner, M. Cummins, and P. Newman. Fast probabilistic labeling of city maps. *Proceedings of Robotics: Science and Systems IV, Zurich, Switzerland*, 2008.

- A. Santamaria-Navarro, E. H. Teniente, M. Morta, and J. Andrade-Cetto. Terrain classification in complex three-dimensional outdoor environments. *Journal of Field Robotics*, 32(1):42– 60, 2015.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.
- B. Suger, B. Steder, and W. Burgard. Traversability analysis for mobile robots in outdoor environments: A semi-supervised learning approach based on 3d-lidar data. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2015.
- G.-Y. Sung, D.-M. Kwak, and J. Lyou. Neural network based terrain classification using wavelet features. *Journal of Intelligent & Robotic Systems*, 59(3):269–281, 2010.
- I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In Advances in neural information processing systems, pages 3104–3112, 2014.
- J. Thiemann, N. Ito, and E. Vincent. The diverse environments multi-channel acoustic noise database (demand): A database of multichannel environmental noise recordings. In 21st International Congress on Acoustics, 2013.
- S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, K. Lau, C. Oakley, M. Palatucci, V. Pratt, P. Stang, S. Strohband, C. Dupont, L.-E. Jendrossek, C. Koelen, C. Markey, C. Rummel, J. van Niekerk, E. Jensen, P. Alessandrini, G. Bradski, B. Davies, S. Ettinger, A. Kaehler, A. Nefian, and P. Mahoney. Stanley: The robot that won the darpa grand challenge. *Journal of Field Robotics*, 23(9):661–692, 2006.
- E. Trautmann and L. Ray. Mobility characterization for autonomous mobile robots using machine learning. *Autonomous Robots*, 30(4):369–383, 2011.
- G. Tzanetakis and P. Cook. Musical genre classification of audio signals. *IEEE Transactions on Speech and Audio Processing*, 10(5):293–302, Jul 2002.
- A. Valada, L. Spinello, and W. Burgard. Deep feature learning for acoustics-based terrain classification. In *Proceedings of the International Symposium on Robotics Research*, 2015.
- B. Verma. Pattern Recognition Technologies and Applications: Recent Advances: Recent Advances. IGI Global, 2008.
- C. Weiss, H. Frohlich, and A. Zell. Vibration-based terrain classification using support vector machines. In *Proceedings* of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 4429–4434, Oct 2006.
- M. C. Wellman, N. Srour, and D. B. Hillis. Feature extraction and fusion of acoustic and seismic sensors for target identification. *Proc. SPIE*, 3081:139–145, 1997.
- W. Zaremba and I. Sutskever. Learning to execute. arxiv preprint arxiv: 1410.4615, 2014.