Semi-supervised Gated Recurrent Neural Networks for Robotic Terrain Classification

Ahmadreza Ahmadi[†], Tønnes Nygaard[‡], Navinda Kottege[†], David Howard[†], Nicolas Hudson[†]

Abstract-Legged robots are popular candidates for missions in challenging terrains due to the wide variety of locomotion strategies they can employ. Terrain classification is a key enabling technology for autonomous legged robots, as it allows the robot to harness their innate flexibility to adapt their behaviour to the demands of their operating environment. In this paper, we show how highly capable machine learning techniques, namely gated recurrent neural networks, allow our target legged robot to correctly classify the terrain it traverses in both supervised and semi-supervised fashions. Tests on a benchmark data set shows that our time-domain classifiers are well capable of dealing with raw and variablelength data with small amount of labels and perform to a level far exceeding the frequency-domain classifiers. The classification results on our own extended data set opens up a range of high-performance behaviours that are specific to those environments. Furthermore, we show how raw unlabelled data is used to improve significantly the classification results in a semi-supervised model.

I. INTRODUCTION

Bio-inspired legged robots offer advantages when walking in extreme environments with their ability to adapt to instantaneous conditions including undulation, slope, roughness, and terrain types. This is possible by changing gaits, foottip arc shapes, footfall placement, stride length, etc., to tune their behaviour and overcome the challenges presented by their environment. Compared to other types of robots, they have more flexibility to effectively couple their hardware and software configuration to the specifics of the terrain.

An important step to fully harnessing these myriad degrees of behavioural freedom is *terrain classification*; the ability for a robot to correctly gauge the type of terrain it is on, and thus enact an appropriate response to overcome the challenges of that terrain. A plethora of previous approaches focus on terrain classification with legged robots using various methods with varying levels of accuracy [1], [2], [3], [4], [5].

In this work, we focus on the use of modern, highly capable deep learning methods, namely *gated recurrent neural networks*, to perform this classification in both supervised and semi-supervised schemes. Our main contributions are:

• A deep learning model for terrain classification via proprioceptive sensing, which significantly outperforms state of the art frequency-domain approaches.



Fig. 1. DyRET quadruped robot on outdoor terrain.

- Comparison of two gated RNN models (LSTMs and GRUs) for terrain classification
- Extensive testing, both on the previously-introduced benchmark PUT dataset [6], [7] which covers indoor terrains, and another large, varied outdoor dataset collected by the authors.
- The first semi-supervised model for robotic terrain classification capable of dealing with raw and variable-length data. We show comparable performance to fully supervised methods, whilst requiring much less anno-tated data.

This work covers three entries of a recent collection of Grand Challenges [8] ¹. Our approach significantly outperforms the frequency-domain models in legged terrain classification [6], and offers the potential for legged robots to generate high-performance behavioural responses that are customised to their environments. Moreover, to the best of our knowledge, this is the first study that investigates a semi-supervised model that can directly use raw and variable-length time-series data for robotic terrain classification. Our dataset² and code³ are available online.

In the remainder of the paper we discuss background research (Section II), introduce our target robot and data sets (Section III), and describe our method and results in Sections IV and V respectively. We conclude with a discussion of our results, and directions for future research in Section VI.

¹Biohybrid and bioinspired robots, Navigation and exploration in extreme environments, and Fundamental aspects of artificial intelligence (AI) for robotics

[†] A. Ahmadi, N. Kottege, D. Howard, and N. Hudson are with the Robotics and Autonomous Systems Group, CSIRO, Pullenvale, QLD 4069, Australia.All correspondence should be addressed to ahmadreza.ahmadi@data61.csiro.au

[‡]T. Nygaard is with the University of Oslo, Norway.

²https://doi.org/10.25919/5f88b9c730442 ³https://github.com/csiro-robotics/ deep-terrain-classification.git

II. RELATED WORK

Legged robot locomotion in rough and unstructured terrain can be addressed either by purely reactive methods or deliberative methods. Reactive methods [9], [10], [11], [12], [13] adapt to new terrain types by changing locomotion parameters or body configuration accordingly. While these methods typically incur low processing overheads, the actual adaptation process can be slow as the robot has to walk over the new terrain for the changes to trigger. Therefore, there is a risk of the robot entering terrain beyond it's locomotion capability before it is identified, potentially stranding the robot in a local minima.

Deliberative methods first classify the type of terrain, then switch to appropriate behaviours rather than adapting reactively. The first step in this process is terrain classification [1], [2], [3], [4], [5]. Proprioceptive sensing and visual perception are often used to gain information about the terrain through the robots' joints, force sensors attached to the feet, inertial measurement units (IMUs), and cameras. These input sequences are then used in a machine learning framework to classify the type of terrain the robot is traversing. Methods based on visual perception may fail in the presence of large variations in illumination intensity, dust, fog, smoke, or seasonal changes in color perception, etc. The state of the art in proprioceptive sensing is dominated by frequency domain classifiers [14], [6], [15]. Variable length input sequences are encountered especially with legged robots, e.g., when the robot walks at different speeds resulting in variable step frequencies. If this is not addressed, the classification results will only be valid when the robot is walking through terrain at the speed used during the training phase. Recent neural approach [6] used an FFT-based pre-processor to overcome problems faced by neural network-based methods when processing variable length sequences.

Recurrent Neural Networks (RNNs) have been widely used in natural language processing tasks as a next step predictor, meaning that they predict what comes next in a sequence [16], [17], [18], [19]. Here, RNNs are used in an unsupervised paradigm as data labelling is not required during training. Semi-supervised learning combines supervised and unsupervised learning to deal with the major drawback of supervised models, an extensive reliance on hand-annotated datasets. Semi-supervised learning has been widely studied for image classification, natural language processing, and video prediction, and achieved outstanding results with small amounts of labelled data [20], [21], [22], [23], [24], [25].

Applications of semi-supervision for robotic terrain classification are underrepresented in the literature. A semisupervised Laplacian Support Vector Machine (SVM) was recently proposed for a relatively small dataset (1584 samples for 6 different types of terrains); it was shown that the proposed method could achieve higher accuracy than a traditional Laplacian SVM [26]. Ten time-domain features were used as input, such as number of sign variations in a sample, sample mean, sample variance, an auto-correlation function of a sample, an impulse factor, etc. In other words, the method does not deal with raw data, and requires many hand crafted features to function. Furthermore, the proposed method cannot deal with variable-length datasets such as PUT and QCAT because the length of data needs to be set in advance to extract time-domain features.

We are therefore motivated to investigate the utilisation of RNNs as semi-supervised models for robotic terrain classification, to deal with raw and variable-length time-series data. We propose a semi-supervised method that incorporates unsupervised RNNs and supervised RNNs, and test on two datasets (one indoor, one outdoor). We demonstrate (i) our time domain representation outperforms recent frequencydomain representation, and (ii) semi-supervision provides equivalent performance to state of the art algorithms, whilst requiring much less labelled data.

III. ROBOTS AND DATASETS

In this section we introduce our robotic test platform, together with the PUT [6] and QCAT datasets used in the experimentation. Details of the two datasets are provided in Table I.

A. PUT Dataset (indoors)

The PUT dataset [6] was collected using a six-legged robot platform equipped with a Force/Torque sensor on one of its legs, sampled at 200 Hz. The robot walked at three different speeds, in six walking directions, on six different artificial indoor terrains (sand, rubber, concrete, artificial grass, wood chipping and gravel). Eighty steps are completed for each combination of speed, direction, and terrain, giving a total of 8640 steps in the whole dataset.

B. Test Platform

We used the open-source Dynamic Robot for Embodied Testing (DyRET) platform [27] (Fig. 1) to collect the data. DyRET is a mammal-inspired robot weighing about 5 kg, built specifically to facilitate machine learning research on real-world robots [28]. Control is via position controlled Robotis servomotors and a high level spline-based gait controller [29]. The robot is equipped with individual directional force sensors (OptoForce OMD-20-SH-80N) on each foot, reporting force on three axes at 100 Hz. An Attitude and Heading Reference System (XSens MTI-30) reports linear acceleration, rotational velocity and orientation at 100 Hz.

C. QCAT Dataset (outdoors)

The QCAT dataset was collected at different locations on CSIRO's QCAT site in Brisbane, Australia, in November 2019. Fig. 2 shows the different environments comprising the data set: a concrete road, grass, gravel, mulch, a dirt path, and sand. Data collection involved walking with a fixed gait, but with three different step frequencies (0.125 Hz, 0.1875 Hz and 0.25 Hz) and two different step lengths (80 mm and 120 mm), for a total of six different speeds tested per surface. The robot walks forwards for eight steps, with ten repeats for a total of 80 steps per speed and surface. To be representative of the terrain type, each repeat occurs on a

TABLE I

STATISTICS FOR THE TWO DATASETS USED IN THIS PAPER. DYRET IS SYMMETRIC WITH SENSORS ON BOTH FRONT AND BACK FEET. THE DATA REPRESENTS BOTH FORWARD AND REVERSE WALKING DIRECTIONS.

	QCAT dataset	PUT dataset
Number of surfaces	6	6
Number of speeds	6	3
Number of directions	*1	6
Sample rate	100 Hz	200 Hz
Number of sensors	4	1
Steps per combination	80	80
Total number of samples	2880	8640
Walking duration	222 min	-

different part of the terrain. The dataset consists of the force senors' measurements (12 dimensions: 4 sensors \times 3) and the IMU sensor's measurements (10 dimensions: 3 of linear accelerations, 3 of angular velocities, and 4 of orientations).

IV. METHOD

Recurrent Neural Networks (RNNs) are a prevalent machine learning technique for dealing with time-series data. An RNN takes an external signal x_t and the previous hidden state h_{t-1} as input, and outputs the current hidden state h_t as:

$$\boldsymbol{h}_t = f(\boldsymbol{W}\boldsymbol{x}_t + \boldsymbol{U}\boldsymbol{h}_{t-1}) \tag{1}$$

where W and U are learnable parameters and f is a non-linear activation function. One advantage of RNNs is their ability to deal with variable-length input sequences. RNNs are able to learn a distribution over a variablelength sequence by learning the distribution over the next input [30]. Early RNN models had difficulties dealing with long-term dependencies in data [31], [32], which led to the development of architectures with more direct memory mechanisms including memory registers, and gated activation functions replacing simple non-linear activation functions. In this paper we focus on two popular implementations; Long Short-Term Memory (LSTM) [33] and Gated Recurrent Unit (GRU) [30].

A. Long Short-Term Memory

An LSTM cell (Fig. 3a) comprises a memory cell c_t , an input gate i_t , a forget gate f_t , and an output gate o_t as:

$$h_{t} = o_{t} \tanh(c_{t})$$

$$o_{t} = \sigma(W_{o}x_{t} + U_{o}h_{t-1})$$

$$c_{t} = f_{t}c_{t-1} + i_{t}\tilde{c}_{t-1}$$

$$\tilde{c}_{t} = \tanh(W_{c}x_{t} + U_{c}h_{t-1})$$

$$f_{t} = \sigma(W_{f}x_{t} + U_{f}h_{t-1})$$

$$i_{t} = \sigma(W_{i}x_{t} + U_{i}h_{t-1})$$
(2)

The input, forget, and output gates control how much new information is memorized, how much old information is forgotten, and how much information is output from the memory cell, respectively.

B. Gated Recurrent Unit

A GRU unit (Fig.3b) consists of a reset gate r_t , and an update gate z_t , which reset and update the memory content adaptively.

$$h_{t} = (1 - z_{t})h_{t-1} + z_{t}\bar{h}_{t}$$

$$z_{t} = \sigma(W_{z}x_{t} + U_{z}h_{t-1})$$

$$\tilde{h}_{t} = \tanh(Wx_{t} + U(r_{t} \odot h_{t-1}))$$

$$r_{t} = \sigma(W_{r}x_{t} + U_{r}h_{t-1})$$
(3)

where \odot denotes element-wise multiplication. GRUs can be considered as an adaptive leaky integrator neuron [34] in which z_t is a user-defined constant value.

C. RNN for Terrain Classification

As shown in previous sections, GRUs have a simpler structure and fewer learnable parameters, and are therefore computationally more efficient compared to LSTM. However, it has been shown that RNN performance critically depends on the task and dataset [35]. We evaluate both GRUs and LSTMs on multiple data sets in this paper. Prior work evaluated LSTMs on the PUT data set, [6], [7] but reports poor results on the raw data, producing predictions not much greater than random chance. Instead they developed methods in the frequency domain [6] with 80.39% accuracy, and fixed data windowing methods using a one dimensional Convolutional Neural Networks (CNNs) input layer [7] fed into LSTMs which achieved 96.89%, at the cost of data windowing. In our work we rigorously test both GRUs and LSTMs on the raw PUT dataset without windowing, training on multiple subsets of the data using k-fold cross-validation and reliably demonstrate high performance results using our RNN implementations provided as part of this paper. Prior work showed results for a single 90/10 train/validation split (without a holdout test data set or using cross-validation) in which the model may be overly biased and overfitted. We show comparable results using k-fold cross-validation (avoiding bias), without additional efforts to transfer or window the data.

We also tested LSTMs and GRUs for the QCAT dataset to see whether gated RNNs can deal with a variable-length outdoor natural environment dataset. Fig. 4(a) shows how an RNN can be used for this task. The model takes an external signal $x_{1:T}$ from time step 1 to T during forward computation and only outputs y at time-step T. The model does not need to output at every time-step because the whole $x_{1:T}$ belongs to one class. This is different from a regression task where a regression model outputs a different signal at each time-step. As both datasets have variablelength sequences, T has a different value for each sequence. The output y is computed by softmax, which produces a distribution over the terrain classes. In addition, we tested semi-supervised RNNs for both datasets to examine their performances when only a small portion of annotated data was provided. The proposed model is shown in Fig. 4(b), in which a classifier model is stacked above a predictor model.



Fig. 2. The different terrains used for collecting our dataset (QCAT). Left to right: Concrete, Grass, Gravel, Mulch, Dirt and Sand.

V. RESULTS

We conducted our experiments to examine how gated RNNs perform in terrain classification tasks for legged robots in both supervised and semi-supervised fashions. The first experiment investigates GRUs and LSTMs for classifying terrains in the (indoor) PUT dataset, and the second experiment does the same on the outdoor QCAT dataset. The last experiment examines semi-supervised RNNs for both datasets. In all experiments, each dimension of the input data is normalized to zero mean and unit standard variation. The cross entropy between the ground truth data \hat{y} and the RNN output y was used for the loss function as:

$$\boldsymbol{L}_{loss} = \sum_{c=1}^{C} \hat{\boldsymbol{y}}_c \log \boldsymbol{y}_c \tag{4}$$

where C is the total class number. An L2 regularization term for N learnable parameters was also added to the loss as

$$\boldsymbol{L}_{loss} = \sum_{c=1}^{C} \hat{\boldsymbol{y}}_c \log \boldsymbol{y}_c + \lambda \sum_{n=1}^{N} \boldsymbol{\theta}^2$$
(5)

As neither data set has an independent set of test samples available, we apply k-fold cross-validation to reduce bias. Following commonly-used settings, we set k to 5 and 10 in our experiments [36], [37].



Fig. 4. RNNs for terrain classification in a (a) Supervised fashion and (b) Semi-supervised fashion. Figure (a) shows an unfolded structure of the RNN model, while figure (b) illustrates the folded RNN models with cyclic representations. The h_t^p and h_t^c show the hidden states of the predictor RNN and the classifier RNN, which are stacked one above the other followed by a Fully Connected Layer (FCL).

A. RNNs Trained on the PUT Data Set

We evaluated GRUs and LSTMs for classifying the terrains in the PUT dataset. The mean and Standard Deviation (SD) results of models trained by a simple loss function (Eq. 4) are shown in Table II. The Minimum (Min) accuracy among k models and Maximum (Max) accuracy among k models are also given. First, the most standard and straightforward architectures of GRUs and LSTMs are compared, then we added more complexity to the superior RNN to achieve higher accuracy. GRUs consistently outperform LSTMs, a trend which is repeated for 10-fold Cross-Validation (CV) models. Ten-fold models are noted to always have a better mean accuracy than 5-fold CV models. However, SD values are smaller for some of the 5-fold CV models. Table III illustrates the results of the models trained by loss with the L2 regularization term (Eq. 5) in which λ was set to 0.01. The performance of GRUs is improved by adding the regularization term, which is not always the case for LSTMs. One interesting performance is that even small GRUs (50 units) achieved a respectable 89.47% accuracy. This can be useful if one wants to use the classifier on embedded systems where memory efficiency is vital. The main result of note is that the more complex architecture (RNNs+FCLs) achieved a mean accuracy of 93.2%, which is significantly better (12.81%) than the model reported in [6], and 26.2% better than the best results obtained by a conventional machine learning algorithm (SVMs), reported in [6]. The more complex architecture consists of GRUs plus Fully Connected Layers (FCLs) with dropouts [38], and learning rate decay was used during its training phase. The notable network achieved 95.39% maximum accuracy, whereas its minimum accuracy is 92.06%. This difference shows the importance of using cross-validation and reporting mean and standard deviation of all CV models.

The overall results show that RNNs are capable of dealing with raw PUT data without reducing the temporal size or transferring from the time-domain to the frequency domain. Our results therefore suggest that gated RNNs are capable of dealing with long sequences specially when datasets are not too complex such as PUT. This also positions GRUs as a leading candidate for learning terrain classifiers in this context. We later show the importance of keeping the temporal resolutions for RNN models (Fig. 6).

Using a straightforward architecture, we show that our RNN implementation performs significantly better than previously reported results given the raw data[6], [7]. For example, 50 LSTMs with no regularization and a fixed learning rate achieved a mean accuracy of 87.68%, compared to 18% in the literature.

TABLE II

Accuracy for PUT dataset. SD, Min, Max, and CV abbreviate standard deviation, minimum accuracy among k RNN models, maximum accuracy among k RNN models, and cross validation, respectively.

		Mean%	SD%	Min%	Max%
	50 GRUs	87.68	1.02	85.84	89.06
	50 LSTMs	83.68	1.06	81.35	84.99
10 fald CV	100 GRUs	90.15	1.11	88.30	91.42
10-1010 C V	100 LSTMs	85.33	1.39	82.64	87.67
	400 GRUs	92.13	0.48	91.42	92.81
	400 LSTMs	90.43	0.59	89.6	91.43
	50 GRUs	82.93	1.13	84.93	88.04
	50 LSTMs	81.86	0.76	80.47	82.73
5-fold CV	100 GRUs	90.2	0.6	89.49	90.94
	100 LSTMs	83.98	0.67	83.21	85.2
	400 GRUs	91.42	0.77	90.24	92.54
	400 LSTMs	89.67	0.74	88.95	91.1

TABLE III

ACCURACY FOR PUT DATASET WITH REGULARIZATION LOSS FOR TRAINING OF NETWORKS.

		Mean%	SD%	Min%	Max%
10-fold CV	50 GRUs	89.47	0.79	88.3	90.77
	50 LSTMs	83.63	1.85	81.22	87.34
	100 GRUs	90.71	0.87	89.7	92.49
	100 LSTMs	85.79	1.64	83.58	88.52
	400 GRUs	92.35	0.78	90.67	93.56
	400 LSTMs	88.3	1.49	85.74	91.0
	RNNs+FCL	93.2	0.89	92.06	95.39
5-fold CV	50 GRUs	88.44	1.06	86.86	90.08
	50 LSTMs	83.07	0.46	82.31	83.6
	100 GRUs	89.99	0.68	88.95	90.88
	100 LSTMs	85.44	0.38	85.09	86.11
	400 GRUs	91.91	0.41	91.26	92.38
	400 LSTMs	85.34	1.95	83.7	89.06

We suggest two possible reasons for poor performance observed in previous results. First, variable length sequences in the PUT dataset use zero-padding, which may perform poorly in certain (static) RNN implementations. Secondly, the length of the sequence may not have been made available to the RNN, which forces common RNN implementations to revert from dynamic to static, causing the problem identified above.

Specifically, padded zeros affect both forward computation and back-propagation through time during training. In Eq. 1, the value of current hidden states h_t depends on values of both input x_t and previous hidden states h_{t-1} , therefore zero values of input x_t will not necessarily result in zero values of current hidden states h_t . This means padded zeros can deteriorate the performance of RNNs if not handled properly.

B. RNNs Trained on the QCAT Data Set

We evaluated GRUs and LSTMs for classifying the terrains on the outdoor, variable-length QCAT dataset. Models were trained on data collected by the force sensors (OptoForce sensors), and the IMU (XSens). IMUs measure movement and orientation, which is affected by many factors other than the surface it walks on. It is also susceptible

TABLE IV

CLASSIFICATION ACCURACY FOR QCAT DATASET WITH REGULARIZATION LOSS FOR TRAINING OF NETWORKS.

		Mean%	SD%	Min%	Max%
	50 GRUs	87.71	2.28	84.37	92.01
	50 LSTMs	80.56	0.96	78.82	81.94
10 fald CV	100 GRUs	89.58	1.24	88.19	92.36
10-101a C V	100 LSTMs	82.29	1.15	80.56	84.38
	350 GRUs	93.02	1.62	90.97	96.53
	350 LSTMs	85.97	1.64	82.29	87.85
	RNNs+FCL	96.6	0.89	95.49	98.61
	50 GRUs	86.6	1.54	85.07	89.41
	50 LSTMs	81.46	1.64	79.17	83.33
5 fold CV	100 GRUs	87.71	1.03	86.46	88.72
5-fold CV	100 LSTMs	83.82	1.67	80.9	85.42
	350 GRUs	91.28	1.05	90.1	92.88
	350 LSTMs	83.19	1.21	81.25	85.07

to noise, which may make the data harder to work with. Force sensors on the feet are more expensive and mounting them can be mechanically and electrically challenging. They are not available on many platforms due to this, but the advantage is that they are in direct contact with the surface, which may reduce noise.

Table IV illustrate the results achieved by GRUs and LSTMs using four force sensors and trained via regularized loss (Eq. 5). Overall, results follow similar trends to those observed with the PUT dataset. GRUs again outperforms LSTMs, and the best network (96.6% of accuracy) belongs to 10-fold CV models with FCL.

We also consider a robot that relies either on a single force sensor, or solely on the IMU. As discussed earlier, force sensors are expensive and many platforms do not have them, demonstrating our method working with a more accessible, reduced sensory payload therefore has merit.

We achieved a mean accuracy of 87.74% when using only one force sensor. Using a cheap IMU sensor (Table V), we see remarkably high mean accuracy (96.63%) that is slightly better than the one achieved by the four force sensors (96.6%). The results suggest that GRUs are capable of dealing with noise in the input signal, showing that

TABLE V ACCURACY FOR QCAT DATASET WITH IMU SENSOR WITH REGULARIZATION LOSS FOR TRAINING OF NETWORKS..

		Mean%	SD%	Min%	Max%
	50 GRUs	90.07	1.77	87.5	93.06
	50 LSTMs	83.06	2.46	78.82	87.5
10 fald CV	100 GRUs	93.06	1.27	91.32	95.49
10-1010 C V	100 LSTMs	84.44	2.52	78.47	88.19
	350 GRUs	94.37	1.68	90.97	96.88
	350 LSTMs	85.0	2.11	81.25	88.19
	RNNs+FCL	96.63	1.17	94.44	98.26
	50 GRUs	88.61	1.11	87.15	90.28
	50 LSTMs	80.83	1.43	78.47	82.64
5 fold CV	100 GRUs	90.63	1.29	88.19	92.01
5-101d UV	100 LSTMs	83.23	1.02	81.6	84.55
	350 GRUs	92.7	0.87	91.32	93.92
	350 LSTMs	84.03	2.17	82.12	87.33



Fig. 5. Comparison of test results between supervised learning and two types of semi-supervised learning. The arrows show the standard deviations. The semi-supervised algorithms significantly outperform the supervised algorithm for both datasets, specially when the percentage of the label data is small. The semi-supervised learning with feature extraction and fine tuning has a superior performance than the semi-supervised one with only feature extraction.

our classification models are widely applicable across many robot platforms. Trends are consistent with previous results obtained using force sensors. Notably, more of the 5-fold CV models had less variance than the 10-fold CV models. We also investigated the models given both IMU and force inputs and observed that they could achieve the accuracy of $97.64 \pm 1.03\%$. Given the remarkable accuracy obtained by IMU or force sensors, it does not seem necessary to have both types of sensors for robotic terrain classification, although more investigations are required to ascertain the limits of this, or applicability to different types of robots.

C. Semi-supervised learning

As shown for both PUT and QCAT datasets, supervised RNNs achieved high accuracy for classifying the different terrains. However, the process of data collection for supervised models is tedious. For data annotation, either the robots need to walk on different types of terrains separately or one needs to hand-label the data later on. The effort can be reduced if only a portion of data needs to be annotated for comparable performance. To that end, we proposed semisupervised RNNs for the terrain classification to investigate the possibility of reducing hand-labeling efforts.

In the proposed semi-supervised model, we stacked un-

supervised RNNs, supervised RNNs, and Fully Connected Layer (FCL) neural networks (Fig. 4(b)). First, the unsupervised RNNs, also referred as the predictor RNNs, are trained by taking the current signal x_t (Force/Torque signals for PUT and IMU signals for QCAT) as input and predicting the next step signal x_{t+1} in their outputs. A portion of input data (predicting data) was used for training of the predictor RNNs. This method of training is called unsupervised learning because the target signals are provided by input signals themselves and no human efforts for data annotations are needed. After training, the weights are frozen and the supervised RNNs, also referred as the classifier RNNs, and the FCL are trained by using other portion of the data (classifying data). The classifier models take the predictor RNN outputs as input and estimating the outputs y, types of terrains, at time-step T. We refer to this model as Feature Extracting (FE) semisupervised learning because the predictor RNNs were used as feature extractor. We also investigated fine-tuning of the whole FE semi-supervised learning model by decreasing the learning rate and retraining the whole model with classifying data. This model is referred as Feature Extracting (FE) + Fine Tuning (FT) semi-supervised learning model. Both models are compared on PUT and QCAT datasets. For sake of space and given that GRUs outperformed LSTMs in all previous experiments, only GRUs are used in the following experiments. IMU data is only used for the QCAT dataset given our previous promising results with this cheap, available sensor.

We investigated 5 different splits of predicting/classifying data. In Fig. 5, 5% means that the whole dataset was randomly divided into two sections: 90% of data was used as the predicting data and the remainder was classifying data. Using 2-fold cross-validation to lower classifier bias, the classifying data (10% of the whole data) was randomly split into 2 parts. Therefore, the supervised model (classifier RNNs and FCL) was only trained by 5% of the whole data, and evaluated on the other 5%. As 5% of the data was used for training of the supervised model, we called these models '5%'. We tested the classifiers on the rest 90% of the whole data, which are referred as test accuracy. Similarly, '25%' means that the whole dataset was randomly divided into two sections: 50% predicting data and 50% classifying data. All networks were trained by 25% of the data, evaluated on 25%, and tested on 50%. In order to compare the semi-supervised models with supervised models and show the effect of the predictor RNNs in the semi-supervised models, we trained supervised classifiers using the classifying data.

Test accuracy results (Fig. 5) shows the superior performance of semi-supervised models over supervised models. The gap between their percentage accuracy is significantly larger for a smaller amount of labels, and becomes smaller as more labels are provided during training. Supervised model accuracy are relatively low (49.31% and 61.85%) when networks are trained by 5% and 10% labels, possibly because QCAT is a relatively small dataset. The effect is less evident for PUT, which is larger.

Results show the effectiveness of fine-tuning for semi-



Fig. 6. Principle component analysis of GRU hidden states given 720 testing samples at different time-steps. The points belonging to the same class have clustered together as time evolves. Some classes (sand and concrete) seem to be easier for the RNN model to classify than others, specifically mulch, dirt, and gravel.

supervised learning, as those models outperform the FE semi-supervised models in all cases. Accuracy is slightly lower for 25% (92.24% accuracy) than 20% (91.86% accuracy) for PUT dataset, and it is really close for QCAT. This is likely because predictor RNNs are trained on less data (50%) for the former model, although its classifier model had access to 5% more labels, suggesting that a balance between amounts of training data for the supervised and unsupervised model may be needed to achieve the highest possible accuracy using semi-supervised learning is an effective learning method when few labels are available, addressing a major shortcoming of supervised learning, especially for terrain classification where extensive annotations may be difficult to procure .

We visualised the temporal evolution of hidden states of a GRU given 720 testing samples of different lengths (T). Principal Component Analysis (PCA) reduced the GRU hidden state dimension (200 GRU units) to two principle components. Fig. 6 displays the PCA results at time-step t= 10%T, 40%T, 70%T and T for QCAT dataset. The dirt, mulch, and gravel classes are seen to be the most challenging ones for the network due to similarities between these three classes (Fig. 2). Sand and concrete are the most separable classes for the network due to their contrasting mechanical properties. The figure also shows how the hidden states of the same classes are mostly clustered together as time goes by. This emphasizes the importance of the data temporal resolution for the terrain classification task.

We investigated other possible architectures for semisupervised learning, which have been shown to be effective in tasks such as image classification or NLP. Our first approach replaced predictor RNNs with fully connected layers. Second, we removed the predictor RNNs, pre-trained RNNs and fully connected layers for the one-step prediction task, and fine-tune the whole network for the classification task. We also replaced the predictor RNNs with auto-encoder RNNs, meaning that the unsupervised RNNs were trained on the current signal x_t (Force/Torque signals for PUT and IMU signals for QCAT) and predicted the the same time step signal x_t . Preliminary results suggest that non of the aforementioned semi-supervised methods are effective for the robotic terrain classifications (PUT and QCAT), they under-performed the supervised models, although more detailed analysis and investigations are needed to have a conclusive statement.

VI. CONCLUSIONS

We utilized supervised and semi-supervised gated RNNs for the robotic terrain classification. Our classifiers were first evaluated on the PUT dataset composed of time-series data with variable lengths that were collected in an indoor environment. RNN models given time-series dataset significantly exceeded the accuracy rates of the SVM and the fully connected neural model using frequency-domain transferred data. Furthermore, we achieved high accuracy rates by RNN classifiers on our own dataset composed of timeseries data with variable lengths that were collected in an outdoor environment. The results obtained from both datasets suggest that a GRU outperforms LSTM for proprioceptive terrain classification. The results also show the importance of the data temporal resolution for terrain classification. In the second experiment, we showed that IMU sensors, available on many robot platforms, may be a sensor of choice for terrain classification. We introduced the first deep semi-supervised models for robotic terrain classification, and showed that they are capable of directly dealing with raw and variable-length time-series data. Results indicate that semi-supervised models outperformed supervised models remarkably when only small amounts of annotated data are available, suggesting that less annotated data is required for terrain classification, and thus larger usable data sets can be easily made available. The results open up an interesting future extension of the current work: a transfer learning between robot platforms, meaning that a different robot will be used for training the unsupervised model with no annotated data. Another interesting future work can be to investigate probabilistic RNNs as semi-supervised models for robotic terrain classification. Probabilistic RNNs [39], [40], [41] are known to attain better generalization capabilities than deterministic ones specially for dealing with data with high variability as evidenced in this domain.

ACKNOWLEDGEMENTS

We would like to give our special thanks to people who helped us with the current study. we are particularly grateful for great assistant given by Benjamin Tam for data collection.

REFERENCES

- M. A. Hoepflinger, C. D. Remy, M. Hutter, L. Spinello, and R. Siegwart, "Haptic terrain classification for legged robots," in *IEEE International Conference on Robotics and Automation*, 2010, pp. 2828–2833.
- [2] P. Filitchkin and K. Byl, "Feature-based terrain classification for littledog," in *IEEE/RSJ International Conference on Intelligent Robots* and Systems, 2012, pp. 1387–1392.
- [3] G. Best, P. Moghadam, N. Kottege, and L. Kleeman, "Terrain classification using a hexapod robot," in *Australasian Conference on Robotics* and Automation, 2013.
- [4] J. Christie and N. Kottege, "Acoustics based terrain classification for legged robots," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 3596–3603.
- [5] X. A. Wu, T. M. Huh, R. Mukherjee, and M. Cutkosky, "Integrated ground reaction force sensing and terrain classification for small legged robots," *Robotics and Automation Letters*, vol. 1, no. 2, pp. 1125–1132, 2016.
- [6] J. Bednarek, M. Bednarek, L. Wellhausen, M. Hutter, and K. Walas, "What am i touching? learning to classify terrain via haptic sensing," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 7187–7193.
- [7] J. Bednarek, M. Bednarek, P. Kicki, and K. Walas, "Robotic touch: Classification of materials for manipulation and walking," in *IEEE International Conference on Soft Robotics*, 2019, pp. 527–533.
- [8] G.-Z. Yang, J. Bellingham, P. E. Dupont, P. Fischer, L. Floridi, R. Full, N. Jacobstein, V. Kumar, M. McNutt, R. Merrifield, B. J. Nelson, B. Scassellati, M. Taddeo, R. Taylor, M. Veloso, Z. L. Wang, and R. Wood, "The grand challenges of science robotics," *Science Robotics*, vol. 3, no. 14, 2018.
- [9] N. Kottege, C. Parkinson, P. Moghadam, A. Elfes, and S. P. Singh, "Energetics-informed hexapod gait transitions across terrains," in *IEEE International Conference on Robotics and Automation*, 2015.
- [10] M. Bjelonic, N. Kottege, and P. Beckerle, "Proprioceptive control of an over-actuated hexapod robot in unstructured terrain," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 2042–2049.
- [11] D. Williamson, N. Kottege, and P. Moghadam, "Terrain characterisation and gait adaptation by a hexapod robot," in *Australasian* conference on Robotics and Automation, 2016.
- [12] T. Homberger, M. Bjelonic, N. Kottege, and P. V. K. Borges, "Terraindependent motion adaptation for hexapod robots," in *International Symposium on Experimental Robotics*, 2016.
- [13] C. Dario Bellicoso, C. Gehring, J. Hwangbo, P. Fankhauser, and M. Hutter, "Perception-less terrain adaptation through whole body control and hierarchical optimization," in *IEEE International Conference* on Humanoid Robots, 2016, pp. 558–564.
- [14] A. Valada, L. Spinello, and W. Burgard, "Deep feature learning for acoustics-based terrain classification," in *Robotics Research*. Springer, 2018, pp. 21–37.
- [15] C. Bai, J. Guo, L. Guo, and J. Song, "Deep multi-layer perception based terrain classification for planetary exploration rovers," *Sensors*, vol. 19, no. 14, p. 3102, 2019.
- [16] I. Sutskever, J. Martens, and G. E. Hinton, "Generating text with recurrent neural networks," in *ICML*, 2011.
- [17] A. Graves, "Generating sequences with recurrent neural networks," arXiv preprint arXiv:1308.0850, 2013.
- [18] W. Zaremba, I. Sutskever, and O. Vinyals, "Recurrent neural network regularization," *arXiv preprint arXiv:1409.2329*, 2014.
- [19] J. Evermann, J.-R. Rehse, and P. Fettke, "Predicting process behaviour using deep learning," *Decision Support Systems*, vol. 100, pp. 129– 140, 2017.

- [20] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling, "Semisupervised learning with deep generative models," in *Advances in neural information processing systems*, 2014, pp. 3581–3589.
- [21] A. M. Dai and Q. V. Le, "Semi-supervised sequence learning," in Advances in neural information processing systems, 2015, pp. 3079– 3087.
- [22] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko, "Semi-supervised learning with ladder networks," in Advances in neural information processing systems, 2015, pp. 3546–3554.
- [23] D. Maltoni and V. Lomonaco, "Semi-supervised tuning from temporal coherence," in *International Conference on Pattern Recognition*, 2016, pp. 2509–2514.
- [24] W. Lotter, G. Kreiman, and D. Cox, "Deep predictive coding networks for video prediction and unsupervised learning," *arXiv preprint arXiv:1605.08104*, 2016.
- [25] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, "Mixmatch: A holistic approach to semi-supervised learning," in *Advances in Neural Information Processing Systems*, 2019, pp. 5049–5059.
- [26] W. Shi, Z. Li, W. Lv, Y. Wu, J. Chang, and X. Li, "Laplacian support vector machine for vibration-based robotic terrain classification," *Electronics*, vol. 9, no. 3, p. 513, 2020.
- [27] T. F. Nygaard, C. P. Martin, J. Torresen, and K. Glette, "Selfmodifying morphology experiments with DyRET: Dynamic robot for embodied testing," in *IEEE International Conference on Robotics and Automation*, 2019.
- [28] T. F. Nygaard, J. Nordmoen, K. O. Ellefsen, C. P. Martin, J. Tørresen, and K. Glette, "Experiences from real-world evolution with dyret: Dynamic robot for embodied testing," in *Nordic Artificial Intelligence Research and Development*, K. Bach and M. Ruocco, Eds. Springer International Publishing, 2019, pp. 58–68.
- [29] T. F. Nygaard, C. P. Martin, J. Torresen, and K. Glette, "Evolving robots on easy mode: towards a variable complexity controller for quadrupeds," in *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*, 2019, pp. 616–632.
- [30] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," arXiv preprint arXiv:1409.1259, 2014.
- [31] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994.
- [32] Z. C. Lipton, J. Berkowitz, and C. Elkan, "A critical review of recurrent neural networks for sequence learning," *arXiv preprint* arXiv:1506.00019, 2015.
- [33] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [34] W. Gerstner and W. M. Kistler, Spiking neuron models: Single neurons, populations, plasticity. Cambridge university press, 2002.
- [35] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," arXiv preprint arXiv:1412.3555, 2014.
- [36] M. Kuhn and K. Johnson, Applied predictive modeling. Springer, 2013, vol. 26.
- [37] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An introduction to statistical learning*. Springer, 2013, vol. 112.
- [38] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [39] J. Bayer and C. Osendorfer, "Learning stochastic recurrent networks," arXiv preprint arXiv:1411.7610, 2014.
- [40] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, "A recurrent latent variable model for sequential data," in *Advances in neural information processing systems*, 2015, pp. 2980– 2988.
- [41] A. Ahmadi and J. Tani, "A novel predictive-coding-inspired variational RNN model for online prediction and recognition," *Neural computation*, vol. 31, no. 11, pp. 2025–2074, 2019.