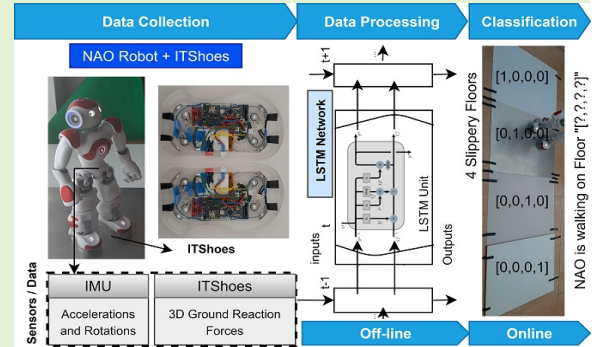

[entry]none/global/

Real-time LSTM-RNN Classification of Floors with Different Friction Coefficients for a Walking Humanoid Robot Wearing a 3D Force System

Luís Almeida, Vítor Santos, and João Ferreira

Abstract—In the study of biped humanoid robots it is crucial to achieve high precision and robustness in locomotion. Humanoid robots that operate in real world environments need to be able to physically recognize different grounds to best adapt their gait without losing their dynamic stability. This work proposes a technique to classify in real time the type of floor from a set of possibilities learnt off-line. Hence, the paper describes the collection and preparation of a dataset of contact forces, obtained with a wearable instrumented system, mixed with the information of the robot internal inertial sensor to classify the type of underlying surface of a walking humanoid robot. For this classification, the data are acquired for four different slippery floors at a rate of 100Hz and it is used as input for a long short-term memory (LSTM) recurrent neural network (RNN). After testing different learning models architectures and tuning the models parameters, a good mapping between inputs and targets is achieved with a test classification accuracy greater than 92%. A real time experiment is presented to demonstrate the suitability of the proposed approach for the multi-classification problem addressed.

Index Terms—Computational Intelligence Techniques, Floor classification, Humanoid Robot locomotion, Wearable Instrumented System, LSTM



I. INTRODUCTION

The study of humanoids increasingly contributes to several scientific areas of which it is possible to extrapolate improvements to our lives, such as in walking rehabilitation, dangerous works, or elderly assistance. To achieve a robot capable of servicing and assisting people, first it has to be able to perform fundamental locomotion tasks such as balancing and walking [pratt2010toward]. Despite the progress and efforts made in the past years, humanoid locomotion is still a challenging problem without definitive solution. Developing a good system that allows a biped robot to walk on unknown and diversified floors, e.g. slippery floors, requires the system to be intelligent and autonomous to adapt in real time so that the robot can successfully overcome the barriers found during locomotion tasks [siciliano2016springer]. In most of the bipedal locomotion approaches, hard contacts with the ground are assumed, although, in real life scenarios this is

normally not accurate. Despite the several developments over the years, there are no explicit implementations which deal with the changing floor properties [hopkins2015design]. This oversight may lead to disastrous consequences, e.g. the biped robot falls while walking on a non-modeled ground, most of the times preventing the robot to continue its locomotion tasks.

Having an intelligent algorithm that allows the robot to identify the terrain with good accuracy, using force sensors installed on its feet or assembled under it, will give the possibility to eliminate most of the falls while walking on different surfaces [walas2016terrain, multi-axis-review]. The increasing progresses made in the areas of artificial intelligence and machine learning, lead to a significant impact in this field. With the rise of techniques such as neural networks, recurrent neural networks (RNNs), deep learning and reinforcement learning, humanoids can now perform tasks that previously seemed far-fetched.

In this work, the capabilities of a Long Short-term Memory Recurrent Neural Network (LSTM-RNN) (first presented in [schmidhuber1997long]) were analysed to classify the underlying surface of a walking robot. To feed the network, a wearable instrumented system assembled to a robot foot was used to measure the ground reaction forces (GRFs), and the internal robot inertial sensors were used to measure the body accelerations and inclinations on different slippery floors. The

The Fundação para a Ciência e Tecnologia (FCT) is gratefully acknowledged for funding this work with the grant SFRH/BD/136680/2018.

Luís Almeida is with the Institute of Electronics and Informatics Engineering of Aveiro (IEETA), Portugal (luis93.pa@gmail.com).

Vítor Santos is with the Institute of Electronics and Informatics Engineering of Aveiro (IEETA) and with the Department of Mechanical Engineering, Aveiro, Portugal (vitor@ua.pt).

João Ferreira is with the Department of Electrical Engineering, Superior Institute of Coimbra and with the Institute of Systems and Robotics (ISR), Coimbra, Portugal (ferreira@isec.pt).

choice of this type of network was mainly due to its feedback loop which serves as a kind of memory. This means that the past inputs leave a footprint on the model that is expected to be an asset, e.g., when the humanoid robot moves from one floor to another. These networks are capable of recognizing temporal encrypted patterns from dynamic data, which is what happens between the interaction of the humanoid's foot and the different floors; indeed, humanoid walking can be considered a time-dependent task. Additionally, our approach does not require a time window for offline processing, hence, the classification can be done online at every new robot step. The main contribution of our work is the classification of slippery floors using a novel instrumented system that can be adapted to different humanoid robots. This classification will provide knowledge of the characteristics of different floors on which the robot walks to the humanoid controller, thus allowing it to adapt the robot gait according to the different slippery floors requirements.

Several published works in the field of humanoid robots show the applicability and capabilities of the LSTM-RNN [Yu20201335, Ko20201200, Lobos-Tsunekawa20183247, Kerzel2019, Bhattacharjee2018, chalvatzaki2019lstm, bi'lstm]. For example, in [how2014multiple] a LSTM-RNN model was used to classify six different robot behaviours based on ten robot joint time sequences. In [zhao2018collaborative] it was used to generate a robotic motion from the observations of the human movements to achieve fast and responsive human robot collaborative tasks, avoiding the trouble of solving an inverse kinematics or motion planning problem. In [yu2020walking] kinect sensors were used to obtain walking information of the human body under different slopes, and the data collected are used to feed a LSTM neural network to learn the degrees of freedom of multiple lower limb joints, classify and recognize the different slopes and to compensate the robot's ankle joints based on the slope inclination. The results show a robot NAO able to walk on different slopes. In another example [li2018mobile], the authors used a LSTM network to classify motor fault in mobile robots achieving an accuracy of 87%.

In [li'sequential'2021] the authors used the robot's foot soles pressure sensors and inertial measurement unit (IMU) to feed a LSTM network. They used the sensors to calculate two-directions center of pressure (CoP) and the IMU to obtain two-direction acceleration. The authors report that the robot walking process based on the LSTM output is better than when using a fixed gait. Another example using pressure sensors data and a LSTM network is addressed in [zhang'fingervision'2018]. The authors used these sensors to collect data while interacting with different daily objects. The developed LSTM presented a 97.62% accuracy on the test dataset, being that the authors expect that this good slippery classification can improve the robot grasping skills, leading to a better contact and smoother interaction/manipulation.

Different LSTM architectures and their combination with other networks have also been increasingly studied and applied in the humanoid field. In [9292620] the authors used a bidirectional LSTM-based network which makes use of historical measurements of system states to predict

humanoid fall probability in real-time. In [farazi2017online] a deep LSTM network was trained with simulated data and fine-tuned on a set of real data to track and identify Robots with identical appearance. Since the analysed data are transmitted via Wi-Fi, some delay or even data loss are expected, hence the importance of using a LSTM. Similarly, in [zhen'walking'2019] a deep network LSTM was also used, but here the authors fed the network with acceleration signals and used it to detect gait-phases of a walking human, presenting a F-score higher than 92%. In [li'cnn'2019] a LSTM network was combined with a convolutional neural network (CNN) to improve the human-robot interaction. The CNN was used to extract visual features and the LSTM network was used to find the relationship between these features and six basic emotions. The humanoid is able to adapt its response based on the human emotion classified with the LSTM model.

Our approach is comparable to the floor classification problem addressed by [wang2020situ] since a LSTM network is explored to solve the problem, but the approaches differ because [wang2020situ] uses the humanoid foot sinking state and the force reads from a load cell embedded in the robot's ankle to classify deformable terrains, whereas we combine data from the robot inertial sensor and eight force sensors to classify different slippery floors. They achieved 95% accuracy on average during experiments.

The remainder of this work is divided as follows. Section II presents the materials and methods for the data collection and manipulation. The LSTM network implementation, tuning and online experimental results are presented in Section III. Lastly, Section IV presents the conclusions and future challenges.

II. MATERIALS AND METHODS

In previous research activities [almeida2018novel], an instrumented system was developed to be seamlessly assembled on the walking humanoid robot NAO to measure real-time vertical and horizontal ground reaction forces (GRFs). The GRFs are divided into total normal force (vGRFs) and total horizontal force (hGRFs). The developed system is a cost-effective, lightweight and wirelessly instrumented shoe (ITshoe). The ITshoe used on this work is presented in Fig. 1.

The hGRFs can be represented in the sagittal and transverse plane, as depicted in Fig. 1, and calculated as follows:

$$Fh_{st} = \frac{\sqrt{2}}{2} \cdot [(S1_{45} + S2_{45}) - (S3_{45} + S4_{45})] \quad (1)$$

$$Fh_{tt} = \frac{\sqrt{2}}{2} \cdot [(S2_{45} + S4_{45}) - (S1_{45} + S3_{45})], \quad (2)$$

where Fh_{st} is the total horizontal force in the sagittal plane, Fh_{tt} is the total horizontal force in the transverse plane and $S1_{45}$ to $S4_{45}$ are the four Flexiforce sensors, positioned at 45°, used to measure the tangential forces.

The ITshoes are divided into three units (see Fig. 1):

- Sensing unit** - Composed of eight piezo-resistive A301 flexiforce sensors;
- Acquisition unit** - Deals with the electrical conditioning and power supply;




Fig. 1. ITshoe schematic structure. The red block (A) is the acquisition unit, the green block (B) is the streaming unit, and the blue block (C) is the sensing unit. At the top, the image shows the position of the eight force sensors and the reference axis used to decompose the tangential forces.

Streaming unit - Receives the data across a serial communication with the micro-controller and forwards it to the server, through the ESP8266 Wi-Fi module.

The robot NAO also communicates with the server through Wi-Fi and its main software is used to make it walk. The high-level functions allow to define the walking distance and gait configuration parameters, such as, step length, height, frequency and torso translation along the X and Y axis. The server runs the open-source system ROS that deals with all the data flow and communications with the ITshoes and the Robot (for further details, see [almeida2018novel]). The dataset used for this work were recorded using the ITshoes sensors and the humanoid robot NAO internal sensors at a frequency of 100 Hz.

A. Floor multi-classification problem

This work extends the study presented in [almeida2020learning] to classify different slippery floors on which the humanoid robot NAO has to walk while wearing the ITshoes. Our hypothesis is that each floor has a particular characteristic when identified by the ITshoes force sensors together with the robot NAO inertial unit, which is positioned in the center of the body of the humanoid NAO, namely data from its 3-axis gyroscope, acceleration and body inclination. Fig. 2 shows the layout used to collect data for the addressed multi-classification problem.

The system NAO+ITshoes is used to collect data while the humanoid NAO walks on four different slippery floors: Polytetrafluoroethylene (PTFE); Aluminium; Polyethylene high-density (PE-HD); and Melamine floor. We only focused on




Fig. 2. Layout for data collection with four different floors.

the slippery characteristic of these floors since the forces that the humanoid NAO exert on these rigid and flat floors while walking are not significant and do not produce meaningful elastic or permanent deformation. However, our classification can be associated with different material properties, and can also be used to classify different floors such as deformable floors. Table I presents the coefficient of friction (CoF) for the different materials used as the walking floors. Acrylic is the base material of the ITshoes.

TABLE I
FLOOR-SHOE COEFFICIENT OF FRICTION

Material	Coefficient of Friction
Acrylic-Polytetrafluoroethylene (PTFE)	0.11
Acrylic-Aluminium	0.20
Acrylic-Polyethylene high-density (PE-HD)	0.26
Acrylic-Melamine	0.33

B. Data manipulation

Before classifying the floors using the LSTM-RNN, it is necessary to pre-process the raw data. The following procedure describes the developed algorithmic methodology to extract the data that corresponds to the humanoid robot steps, and to format its representation to be used as inputs for the learning approach. The algorithm reads the recorded raw data from the database and outputs only the data that corresponds to moments when the robot's foot is in contact with the floor. The main steps of the algorithm are as follows:

- Use the calibration curves to convert each sensor ($S(i)$) raw data into forces, as given by:

$$F(i) = -\frac{b(i)}{R(i) \left(\frac{1023}{S(i)} - 1 \right)} \frac{1}{m(i)}, \quad i \in \{1, \dots, 8\}, \quad (3)$$

where $R(i)$ is the voltage divider resistor, $m(i)$ and $b(i)$ are the calibration curve slope and y-interception respectively, and 1024 (2^{10}) refers to a 10-bit analog to digital converter (ADC);

- Obtain the start and end of each step: the points where $F_n(i) \approx 0$ (normal force ≈ 0 , robot foot is in the air);
- Use the indices i for which $F_n(i) \neq 0$ to filter the data points for all the recorded variables;
- Normalize each data point to be in the range $[-1, 1]$;
- Reshape the data according to the LSTM model needs.

C. Data for LSTM training, validation and testing

The resulting dataset for this work includes 27800 labeled sequences. Each input sequence consists of 11 features, each composed of 50 samples. Despite the fact that the measured vertical forces are crucial for the data collection and manipulation, these were not used as features to train the network due to the dynamics involved in the robot step being so small that the variability of the normal force sensors are of limited use [almeida2020learning]. The 11 selected features to represent the robot's behaviour on the different studied floors are as follows:

Fh_{st}	Horizontal force in the sagittal plane (N);
Fh_{tt}	Horizontal force in the transverse plane (N);
G_x	Gyroscope X axis (rad/s);
G_y	Gyroscope Y axis (rad/s);
G_z	Gyroscope Z axis (rad/s);
Acc_x	Accelerometer X axis (m/s^2);
Acc_y	Accelerometer Y axis (m/s^2);
Acc_z	Accelerometer Z axis (m/s^2);
BI_x	Body inclination X axis (rad);
BI_y	Body inclination Y axis (rad);
BI_z	Body inclination Z axis (rad).

The data are prepared and randomly divided into three subsets: the training set (60%), which is used for computing the gradient and updating the network weights and biases; the validation set (20%) to measure network generalization and to halt training when generalization stops improving, and the test set (20%) that is used to compare the different LSTM networks, as well as evaluate the ability of the network to correctly classify the floors. To make the classification process more accurate and less biased, a 10-fold cross validation is applied to the data so that the LSTM model is trained 10 times, each time with a different train/test split. Table II shows the final dimensions of the input and target matrices.

TABLE II
DIMENSIONS OF THE INPUT AND TARGET MATRICES

	Input	Target
Training	(16680, 50, 11)	(16680, 1, 4)
Validation	(5560, 50, 11)	(5560, 1, 4)
Testing	(5560, 50, 11)	(5560, 1, 4)

Fig. 3 illustrates the 50 chosen samples of the normal force component from a full humanoid robot step. All the selected features for this classification problem use also 50 samples corresponding to the same time steps.

III. LSTM-RNN EXPERIMENTAL RESULTS

Recurrent Neural networks (RNNs) are the feed-backward version of the conventional feed-forward neural networks. They have a cyclic connection architecture that allows them to update their current state based on past states and current input data. The standard RNN topology suffers from the vanishing gradient [graves'supervised'2012]. To overcome this problem, Schmidhuber and Hochreiter [schmidhuber1997long] developed the Long Short Term Memory (LSTM) unit presented in Fig. 4.




Fig. 3. Illustration of the 50 chosen samples of the normal force of a robot step to be used on the classification of the robot walking floor.

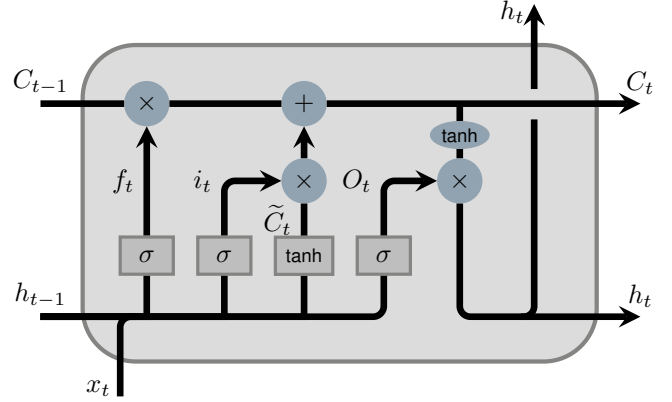


Fig. 4. LSTM cell structure[varsamopoulos2018designing].

The LSTM unit, also known as "memory block" enables the network with the capacity to store and access information over long periods of time. LSTMs achieve this through their 3-gate architecture, which consists of an input, forget and output gates. The input gate decides which input information will be used to update the memory state. The forget gate decides which information to keep or erase from the memory block. Finally, the output gate decides based on the current input and the stored memory which information to output. The LSTM models developed in this work were built using the open source neural network library Keras, running on top of the machine learning platform TensorFlow. The API Keras is written in Python. Several LSTM networks configurations were implemented to classify the different floors from the input layer of 11 features and 50 time steps. The base model used to start training was a LSTM layer with a sigmoid activation function and 50 hidden neurons, a fully-connected output layer with activation function "softmax", and 4 hidden neurons corresponding to the four floors to be classified. The optimizer used in all the training trials was the Adam optimizer [kingma2014adam] because of its ability to converge quickly while traditionally performing better than most other optimizers [chang2018effective, wang2020power]. Additionally, an early stop strategy is applied to the training process to halt it when the loss of the model stops improving. The loss of the model was evaluated using a categorical cross-entropy function, calculated according to (4)

$$L = - \sum_{c=1}^M y_{o,c} \log(p_{o,c}) \quad (4)$$

where M is the number of classes (different floors), $y_{o,c}$ indicates if the class label c is the correct classification for the observation o , and $p_{o,c}$ is the predicted probability of the observation o is of class c . Before tuning the LSTM model hyper parameters, a dropout layer was added to the model since it is commonly used to fight over-fitting and to improve the model performance. This layer is used as a regularization method where input and recurrent connections to LSTM units are probabilistically excluded from activation and weights updates while training a network. It randomly sets input units to 0 with a chosen frequency f_d (fraction of the input units to drop) at each step during training time. Inputs not set to zero are scaled up by $1 \div (1 - f_d)$ such that the sum over all inputs is unchanged. Fig. 5 exemplifies the dropout behaviour for a standard neural network with two hidden layers.




Fig. 5. Dropout behaviour. On the left it is visible a standard neural network and on the right an example of the same network after applying dropout. The units presented with a cross have been dropped [srivastava2014dropout].

Multiple LSTM models architectures were implemented and trained with a varying number of hidden neurons, different activation functions, batch sizes, learning rates and dropout probabilities. Table III shows the evaluated range for some of these parameters as well as the chosen value that produced the overall best results.

TABLE III
RANGE OF NETWORK PARAMETERS EVALUATED.

	Range	Best
Hidden neurons	50 - 500	300
Learning rate	0.00001 - 0.1	0.0001
Batch size	10 - 500	256
Dropout probability	0.1 - 0.5	0.2

Fig. 6 shows the optimized model obtained for this multi-classification problem using the chosen best parameters. This model has three layers: a LSTM layer with 300 hidden neurons and "tanh" activation functions, a dropout 0.2 layer and a fully-connected layer with four hidden neurons and "softmax" activation functions.




Fig. 6. Optimized LSTM model.

Figure 7 presents the optimized model accuracy/epochs. This model presents a 96.31% and 89.63% training and validation accuracy, respectively.




Fig. 7. LSTM model, Accuracy vs Epochs.

TABLE IV
LSTM MODEL CLASSIFICATION REPORT.

	Precision	Recall	F1-score	Samples
Acrylic-PTFE	0.98	1.00	0.99	1620
Acrylic-Aluminium	0.88	0.92	0.90	1510
Acrylic-PE-HD	0.93	0.84	0.89	1290
Acrylic-Melamine	0.93	0.89	0.91	1140

Table IV presents the classification report of this LSTM model where the metrics Precision, Recall and F1-score are used to evaluate the model's capability of correctly classifying the different floors.

Figure 8 presents the confusion matrix obtained the applying the network to the testing set. From there it can be seen that the model was able to correctly classify 92.09% of the test set data (confusion matrix true positives plus true negatives divided by total of samples), which are data never presented to the network before. Overall, the differences between the high scores achieved for the first class compared to the other classes can be justified not only because of the uneven distribution of the four classes, but also because the floor properties, more specifically the CoF of the first class differs more from the remainder. The gap between the CoF for the remainder classes is similar and, as it can be seen, the results for these classes are very balanced.

Output Class	PTFE	Alum.	PE-HD	Melam.
PTFE	100.0% 1620	2.0% 30	4.7% 60	0.0% 0
Alum.	0.0% 0	92.1% 1390	8.5% 110	7.0% 80
PE-HD	0.0% 0	2.6% 40	84.5% 1090	3.5% 40
Melam.	0.0% 0	3.3% 50	2.3% 30	89.5% 1020
	PTFE	Alum.	PE-HD	Melam.

Target Class

Fig. 8. Confusion matrix of the testing dataset.

IV. ONLINE EXPERIMENT

After training and improving the LSTM network, a real-time experiment of the NAO robot walking on different floors, as illustrated in Fig. 9, was carried out to validate the model. The

gait parameters for the walk are as follows: 0.02 m step length; 0.02 m step height; 50% of robot’s maximum step frequency; and 0.1 rad torso rotation around Y.



Fig. 9. Layout for the real-time on-line experiment.

Figure 10 plots the classification results of the LSTM network, and at the top, some snapshots of the experiment

are shown. In this experiment, the robot walked 1.80 meters on the four floors starting on the PTFE floor and ending on the melamine surface. It is important to mention that we considered as being the correct floor on which the robot walks when the front part of the robot foot is at least 0.04 m over that floor. The detailed results show that the network, although not always with 100 % certainty, classified correctly (with a confidence always larger than 50 %) the different floors for the vast majority of the robot’s steps. Fig. 11 presents the confusion matrix for this experiment, where the LSTM model classified correctly 87 of the 90 ($\approx 97\%$) robot steps. Indeed, the network correctly classified all the steps on the PTFE and PE-HD floors, and misclassified only two steps on the aluminium floor and one step on the melamine floor. Overall, it can be seen that the LSTM model is suitable for this multi-classification problem. The model presented requires approximately 541 milliseconds to classify one input in real time (from data collection to obtaining a network output). Since the data are recorded at 100 Hz and we are using the first 50 data points (500 milliseconds) to feed the model, it means that the model classification time (41 milliseconds) represents only 7.5 % of the total time required for this multi-classification problem. Basically, after one robot step we need to wait 41 milliseconds before taking the procedures to optimize the next step. In the future, if we need to decrease the overall required time, we can focus on strategies to reduce the data points used to feed the model. We tested the LSTM network’s ability to classify correctly the different slippery floors using less timesteps and the results are shown in Fig. 12. We observed for example that the network test accuracy exhibits a loss of less than 4 % and 9 % when using 40 and 30 timesteps, respectively. In Fig. 10 it is visible that the network was still able to correctly classify most of the steps during the online experiment, using the 40 and 30 timesteps model, although with less confidence when compared with the 50 timesteps model. Fig. 13 and 14 show the confusion matrices for these two networks with less timesteps. The LSTM model with 40 timesteps classified correctly 86 of 90 robot steps, that is, one more misclassified step when compared to the 50 timesteps model, and the LSTM model using 30 timesteps only classified correctly 80 of 90, seven less than the original 50 timesteps model. In conclusion, it is possible to reduce the

timesteps used to classify the different floors and still have a model able to classify correctly most of the robot’s steps.

In order to assess the relative importance of the variables involved in this process of floor identification, different training strategies were done using different sets of input variables, namely with horizontal forces only, robot own IMU only, and the combination of both. Fig. 15 shows the average test accuracy of the LSTM model when using only the horizontal forces (HFs), only the IMU data and lastly, both together (HFs+IMU) as was used throughout this work. We can already observe that the data coming from the ITShoes has a much larger performance in the classification than the IMU data alone, but we also observed that when we put together the data coming from the ITshoes and the IMU, the classification results further improve about 8 %.

V. CONCLUSIONS

This work addresses a multi-classification problem using a LSTM recurrent neural network model based on the ITshoes force data together with the robot internal inertial sensor data. A dataset of 27800 labelled steps of a walking biped robot over four different floors were collected and used to train, validate and test the multiple learning models. After several attempts to optimize the LSTM base model, and after tuning its parameters, we achieved a network capable of classifying the different floors with an accuracy of approximately 92 %. The developed online experiment also validated the LSTM-RNN approach to this classification problem, as it correctly classified 87 out of 90 robot’s steps. For future work we will develop a thorough study of the impact of the data coming from the ITShoes and the IMU to validate the importance of the wearable instrumented system. Additionally, in the future it is expected to use this classification to optimize the humanoid robot controller, since it will be easier to achieve an efficient and stable humanoid gait if its controller has information about the floor where it walks on. Following this idea, we also expect that this network can be adapted to be used with different classes by analysing how close a new input will be from one of the studied classes. In fact, we observed that even in cases where classification shows smaller confidence in one material, the next best option is, most often, a material with a neighbor friction coefficient. This shows that this solution is able to detect properly the closest friction coefficient from the set of possibilities trained. So, if enough materials are used to represent the required resolution of friction coefficients, the robot will be able to assess with adequate accuracy the interval of values for the friction coefficient where it is walking on, thus allowing the decision of measures to take to improve walking control. Although this classification is focused on the friction property of the materials, we will also consider to study other characteristics such as damping and stiffness not only to have a richer set of parameters to adapt the robot controller but also to broaden the possibility to detect different floors, i.e deformable floors.

When we classify the different floors we know in advance the floor properties/characteristics and the main challenge where more characteristics of the floors can be handy will be in the future when we try to interpolate this known properties with the classification of unknown floors.

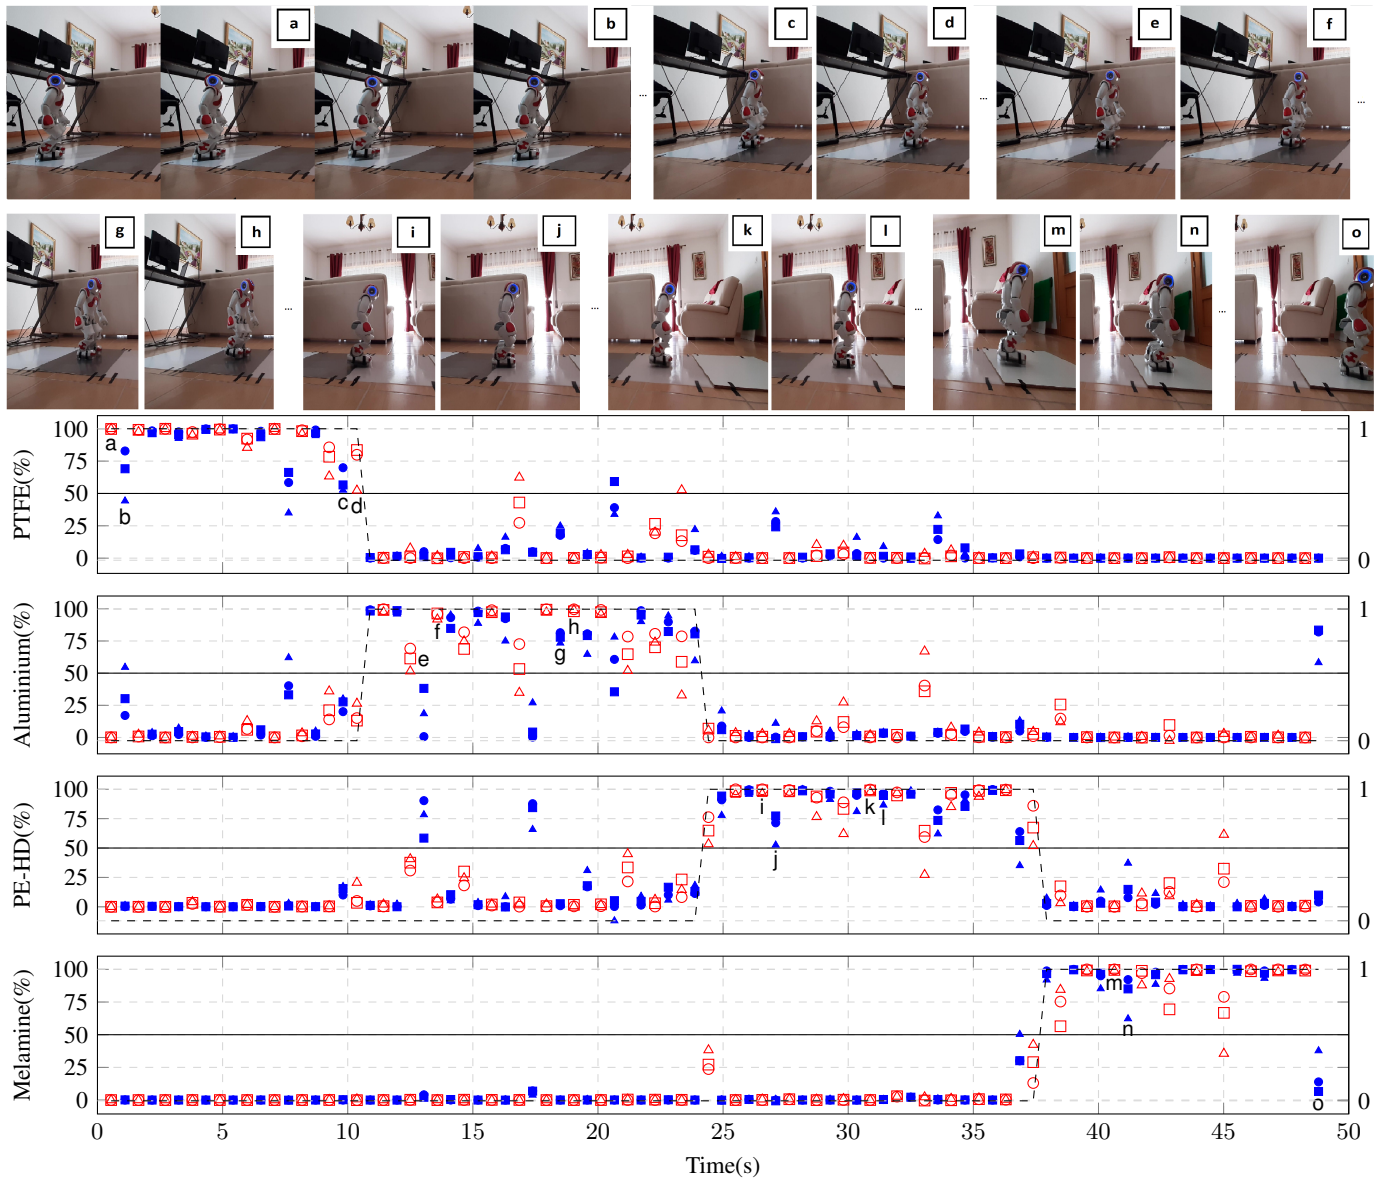


Fig. 10. Individual results on floor classification of NAO walking experiment in real-time using 50, 40 and 30 timesteps. At the top, some snapshots of NAO walking on the different surfaces are presented, and their corresponding output is presented, at the bottom, on the charts for each foot/shoe. \circ Left ITshoe(50); \bullet Right ITshoe(50); \square Left ITshoe(40); \blacksquare Right ITshoe(40); \triangle Left ITshoe(30); \blacktriangle Right ITshoe(30); - - - Real floor on which the robot walks (it is 1 when walking on the current floor).




Fig. 12. LSTM network testing accuracy using different sizes of timesteps.

PTFE	100.0% 19	0.0% 0	0.0% 0	0.0% 0
Alum.	0.0% 0	92.0% 23	0.0% 0	4.8% 1
PE-HD	0.0% 0	8.0% 2	100.0% 25	0.0% 0
Melam.	0.0% 0	0.0% 0	0.0% 0	95.2% 20
	PTFE	Alum.	PE-HD	Melam.
	Target Class			

Fig. 11. Confusion matrix of the real-time experiment.

Output Class	PTFE	100.0% 19	4.0% 1	0.0% 0	0.0% 0
	Alum.	0.0% 0	88.0% 22	0.0% 0	4.8% 1
	PE-HD	0.0% 0	8.0% 2	100.0% 25	0.0% 0
	Melam.	0.0% 0	0.0% 0	0.0% 0	95.2% 20
		PTFE	Alum.	PE-HD	Melam.
		Target Class			

Fig. 13. Confusion matrix of the testing dataset using 40 timesteps.

Output Class	PTFE	89.5% 17	8.0% 2	0.0% 0	0.0% 0
	Alum.	10.5% 2	84.0% 21	4.0% 1	4.8% 1
	PE-HD	0.0% 0	8.0% 2	92.0% 23	4.8% 1
	Melam.	0.0% 0	0.0% 0	4.0% 1	90.5% 19
		PTFE	Alum.	PE-HD	Melam.
		Target Class			

Fig. 14. Confusion matrix of the testing dataset using 30 timesteps.




Fig. 15. LSTM average test accuracy when using only the Horizontal Forces (HF_s), only the data coming from the robot's IMU and all the features (HF_s+IMU) as described in sub-section II-C.




Luís Almeida obtained a 5-year degree in Mechanical Engineering in 2016, at the University of Aveiro, Portugal, where he later, started to pursue his PhD in Mechanical Engineering in 2018. He participated in a research project and some conferences in the area of humanoid robotics. His current research interests include humanoid robots locomotion and control, human-robot interaction and computational intelligent techniques.



Vítor Santos obtained a 5-year degree in Electronics Engineering and Telecommunications in 1989, at the University of Aveiro, Portugal, where he later obtained a PhD in Electrical Engineering in 1995, and the Habilitation in Mechanical Engineering in 2018. He was awarded fellowships to pursue research in mobile robotics during 1990-1994 at the Joint Research Center, Italy. He is currently Associate Professor at the University of Aveiro and lectures courses related to advanced perception and robotics, and has carried out re-

search activity on mobile robotics, advanced perception and humanoid robotics. At the University of Aveiro he has coordinated the ATLAS project for mobile robot competition that achieved 6 first prizes in the annual Autonomous Driving competition and has coordinated the development of ATLASCAR the first real car with autonomous navigation capabilities in Portugal. He is one of the founders of the Portuguese Robotics Open in 2001 and also co-founder of the Portuguese Society of Robotics. His interests extend to humanoid robotics and the application of techniques from perception and mobile robotics to autonomy and safety in Autonomous Driving and Driving Assistance systems.



João P. Ferreira received his BSc degree in Electrical Engineering in 1999, Master in Industrial Automation in 2002 and PhD in Instrumentation and Control in 2010 from the Univ. of Coimbra. Currently, he is a Professor of the Electrical Engineering Department at the Superior Institute of Engineering of Coimbra, Coordinator of specialization course of Industrial Automation, Robotics and Maintenance and Researcher at the Institute of Systems and Robotics of Univ. of Coimbra. He has coordinated and participated in

several funding projects in the area of humanoid and medical rehabilitation robotics, with more than sixty scientific publications in international journals/conferences, over two hundred participations as a reviewer of scientific manuscripts and have a national patent (n^o. 108143). His research interests include humanoid robots, human gait, rehabilitation robotics and artificial intelligence and its application.