Towards intelligent tire development

Tamás Dózsa[1] – Attila Ámon[2] – Ferenc Braun[3] – Ernő Simonyi[4] – Alexandros Soumelidis[5] – János Volk[6] – Péter Kovács[7]

[1,4,5] Institute of Computer Science and Control
[1]dozsa.tamas@sztaki.hu

[4]simonyi.erno@sztaki.hu

[5]alexandros.soumelidis@sztaki.hu

[1,2,7] Eötvös Loránd University
[1]dotuaai@inf.elte.hu

[2]ze3vjn@inf.elte.hu

[7]kovika@inf.elte.hu

[3,6] Nanosensors Laboratory

[3]braun.ferenc@ek-cer.hu

[6]volk.janos@ek-cer.hu

Abstract: In this work we discuss the applications of a piezoresistive 3D force measuring MEMS device with special focus on recent uses in autonomous vehicle control. For such applications, the proposed sensor is implanted into the tires of a vehicle. Because of the rotation of the tires, the produced waveforms display interesting properties which make them suitable for being modelled by so-called adaptive Hermite functions. Such transformations can also be incorporated into neural networks, resulting in simple, interpretable machine learning solutions to heavily nonlinear environmental recognition problems. In this work we review previous applications focusing on the problem of road surface abnormality recognition and provide new, embedded realizations of the above mentioned signal processing algorithms. Our results demonstrate that the proposed measurement system can be used to detect road abnormalities in real time. In addition, we discuss current research trends and potential other applications of the proposed measurement system.

Keywords: Hermite functions; MEMS device; piezoresistive force sensor; smart tires; surface abnormalities; variable projection neural networks, microcontrollers

Introduction

Reliable information on the environment and vehicle dynamics is a key component of any modern vehicle design. This is especially true in applications aiming to achieve fully autonomous behavior of the vehicle. Among the many sophisticated measurement systems developed in recent years [1, 2, 3, 4, 5] to meet the above stated requirements, intelligent tires in particular have found interesting applications [1, 6, 7]. In this work, by „intelligent tires" we are going to refer to any measurement system which uses sensors directly connected to the tires of the vehicle. An example of such a system, which can be found in most new commercial vehicles is active tire pressure monitoring [7] which notifies the driver if the tire pressure falls outside predefined limits. More recently, measurement systems based on intelligent tires enjoyed other applications as well. For example, intelligent tires have been used to estimate the quality of the road surface [1] as well as important markers about vehicle dynamics such as the tire road friction coefficient [6]. These measurement systems employ a variety of different approaches regarding the choice (and installation method) of the sensors and the subsequent signal processing tasks.

In [1], we proposed a novel intelligent tire system built around a piezoresistive force sensor implanted into the tires of the vehicle. This approach allowed us to directly measure tire wall deformations and solved the problem of road surface abnormality detection. Road surface abnormality recognition is an important and well researched environmental recognition task. Precise recognition of obstacles on the road surface, such as potholes and bumps can help to reduce maintenance costs [1, 3]. Furthermore, quickly available, and reliable information on road quality can serve as an input for control systems aimed at modifying vehicle dynamics [8]. Many different approaches have been developed ranging from methods based on image recognition [4, 5] to the use of acceleration sensors [3] and intelligent tires [1]. In our approach using intelligent tires [1], we exploited the special shape of the generated measurements by employing adaptive orthogonal transformations to represent them. Using these transformations, we developed model-driven, explainable machine learning (ML) methods to solve the above-mentioned environmental recognition problem. In this work, we further build upon our results and show that the methods developed in [1] can indeed be used to recognize road abnormalities in real time. Namely, we implement the best performing methods on a microcontroller. In addition, we discuss in detail the benefits of using intelligent tires for environmental recognition, the current direction of our research and the challenges associated with it.

The rest of this manuscript is organized as follows. In section 1, we review our novel measurement system using intelligent tires. We briefly introduce the piezoresistive force sensor employed in our tire design as well as other pieces of hardware required to produce usable measurements. In section 2 we discuss the morphology and specifics of the produced signals. In addition, we review the model driven ML methods required to process the measurements. In section 3, we detail the embedded hardware implementation of our method. Section 4 contains numerical experiments comparing our low-level implementation to the state-of-the-art. We discuss the benefits, current research goals and challenges associated with our measurement system in section 5. Finally, we conclude our results and review possible research directions.

1. Piezoresistive sensor based intelligent tires

In [1], we developed an intelligent tire-based measurement system. The most important hardware component of our system is a 3-D piezoresistive force sensor implanted into the inner wall of the tire. The sensor was installed using the patented method [9], therefore the changes is resistance measured by the force sensor correspond directly to tire wall deformations. It is well known that these deformations of the tire wall are associated with several important markers of the vehicle dynamics [6] as well as road surface quality [1]. The force sensor produces signals along four bridges, whose linear combination can express the force acting on the tip of the sensor using the below relationship:

$$\mathrm{F_x} = \frac{1}{\mathrm{v_0}\alpha_{\mathrm{ls}}\pi_{44}}\left(\Delta\mathrm{V}_{right} - \Delta\mathrm{V_{left}}\right)$$

$$\mathrm{F_y} = \frac{1}{v_0\alpha_{ls}\pi_{44}}\left(\Delta\mathrm{V_{top}} - \Delta\mathrm{V_{bottom}}\right)$$

$$\mathrm{F_z} = \frac{1}{v_0\ \alpha_{ln}\pi_{44}} \times \left(\frac{\Delta\mathrm{V_{right}} + \Delta\mathrm{V_{left}} + \Delta\mathrm{V_{top}} + \Delta\mathrm{V_{bottom}}}{2}\right) \tag{1}$$

In equation (1), $\mathrm{F_x}, \mathrm{F_y}$ and $\mathrm{F_z}$ denote the force components of the tangential and normal directions respectively, while $\mathrm{V_0}$ and $\Delta\ \mathrm{V_{pos}}$ denote common and measured voltages at each bridge. The constants $\pi_{44}, \alpha_{\mathrm{ls}}$ and α_{ln} denote the dominant piezoresistive, linear shear and linear normal coefficients which correspond to the geometric arrangement of the bridges. For a more thorough discussion of the 3-D force sensor and the implantation method, we refer to [1] and [9].It is important to note, that the linear relationship given in equation (1) only holds when we assume that the net force $\mathrm{F} = \mathrm{F_x} + \mathrm{F_y} + \mathrm{F_z}$ acts directly on the sensor. The description of the forces acting between the road and the tire are much more difficult to describe using the force sensor's outputs.

The measurements of the 3-D force sensor are acquired using wireless read out electronics appended to the chassis of the test vehicle above the wheels. The measurements are then scaled and uploaded to the CAN network of the vehicle. The sampling rate used to generate the datasets referenced in this manuscript was 1000 Hz. In our current experimental setup, the above-described measurement system was installed on a Nissan Leaf electric vehicle provided by the Institute of Computer Science and Control of the Eötvös Loránd Research Network in Hungary. For further specification of the system installation and the hardware components of our measurement system we again, refer to [1].

2. Tire sensor signal processing

For the road surface abnormality recognition problem, the data produced by the measurement system described in section 1 was segmented into parts representing full rotations of the tire. The resulting 1-D signals showed quasi compact morphology with a single large peak in the middle which corresponds to the rotation angle, where the force sensor was closest to the ground. In our experiments a road segment was said to be *abnormal* if the vehicle encountered a meaningful obstacle on the road (such as a pothole or bump), while such a 1-D signal representing a single revolution of the tire was recorded. Figure 1. depicts tire revolutions corresponding to normal and abnormal road surfaces. Our main assumption for recognizing road abnormalities based on the produced signals was that signals associated with significant obstacles would contain more noise. This is to be expected, as the shock events created by uneven road surfaces would cause the tire wall to deform quickly which is recorded by the force sensor. Using this hypothesis, we developed several analytical and data driven approaches in [1] to process the signals shown in figure 1. Evaluating the noise level on the tire revolution signals however is not a straightforward problem. This is due to the fact that as the vehicle's speed increases, the effective support of the peaks representing each tire revolution becomes narrower. In turn, traditional methods of measuring noise levels (such as evaluating the empirical standard deviation of each signal) fail to provide a clear distinction between noisy signals and signals acquired at higher speeds. To overcome this problem, we proposed an adaptive high pass filtering scheme in [1], which removes the main peak of each signal while preserving the noise components appearing because of uneven road surfaces intact. This algorithm was realized by approximating each tire revolution signal using adaptive Hermite expansions. Finally, instead of measuring the noise levels on the transformed signals using a fixed statistic, we showed in [1] that data-driven methods can make the abnormality recognizer algorithms more accurate. In particular, we showed that if the parameters of the adaptive orthogonal transformation realized by Hermite functions were trained together with the weights of an underlying fully connected neural network, the resulting model-based abnormality classification scheme could outperform pure data-driven and analytic approaches alike. Below, we give a brief review of our recommended abnormality recognizer algorithm. We note in advance, that the presented model-based neural network follows the so-called variable projection network (VP-NET) architecture introduced in [10].

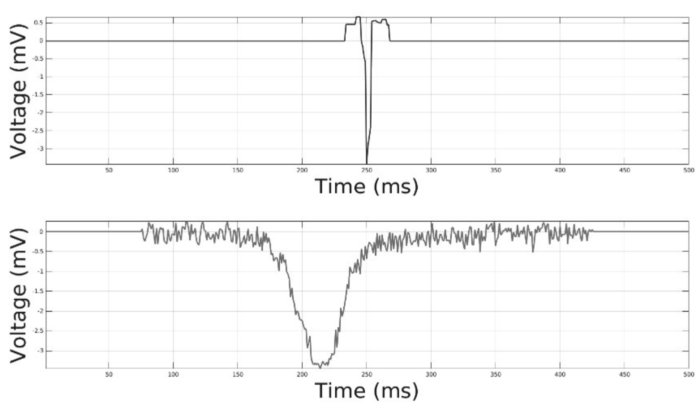

Figure 11. TOP: normal tire revolution signal. BOTTOM: abnormal tire revolution signal. Abnormal signals contain more noise caused by surface abnormalities.

2.1 Adaptive Hermite expansions

As described above, our goal was to identify a mathematical transformation which removes the main peak of the signals corresponding to tire revolutions while keeping the noisy peaks appearing because of road abnormalities intact. To achieve this, we considered the difference between the measured signals and their smooth approximations acquired through adaptive Hermite expansions [1, 10, 11]. Adaptive

Hermite functions, first introduced in [11] have proven to be a useful tool for representing quasi compactly supported signals across a variety of applications. In particular, Hermite functions (and their generalizations) have been successfully applied to process medical signals such as ECG and action potentials [10, 11].

Formally, adaptive Hermite functions can be given as

$$\varphi_k^{\lambda,\tau}(x) \coloneqq \sqrt{\lambda} \cdot h_k\big(\lambda(x-\tau)\big) \cdot e^{-\frac{(\lambda(x-\tau))^2}{2}} \cdot \left(2^n n! \sqrt{\pi}\right)^{-\frac{1}{2}} \tag{2}$$

where h_k denotes the k-th orthogonal Hermite polynomial [12, 13]. Adaptive Hermite functions form a complete and orthogonal function system in $L_2(R)$, therefore any $f \in L_2(R)$ function can be expressed as

$$f = \sum_{k=0}^{\infty} \langle f, {\varphi_k}^{\lambda,\tau} \rangle {\varphi_k}^{\lambda,\tau} \tag{3}$$

where convergence is understood in the norm generated by the usual L_2 inner product denoted by $\langle \cdot,\cdot \rangle$. In (2), parameters λ and τ denote a dilation and translation transformation, respectively. In the case of modeling tire revolutions this is especially useful, because as described above the length of the support of the main peaks appearing in the signals corresponds to the speed of the vehicle. Adaptive Hermite functions are able to model this change with the dilation parameter λ (see figure 2). Hermite functions can be used to approximate many types of signals because their modulus tends quickly to zero as $|x| \to \infty$ and the first few Hermite functions resemble the morphology of waveforms appearing in a variety of problems (e.g. ECG signals, tire revolution signals). This means that in practice truncating the expansion (3) at an index $n \in N$:

$$f \approx \sum_{k=0}^{n-1} \langle f, {\varphi_k}^{\lambda,\tau} \rangle {\varphi_k}^{\lambda,\tau} \tag{4}$$

still provides a good approximation of the signal, even if n was a small number. The approximation given in (4) can be considered the orthogonal projection of the signal f onto the subspace spanned by the first n adaptive Hermite functions.

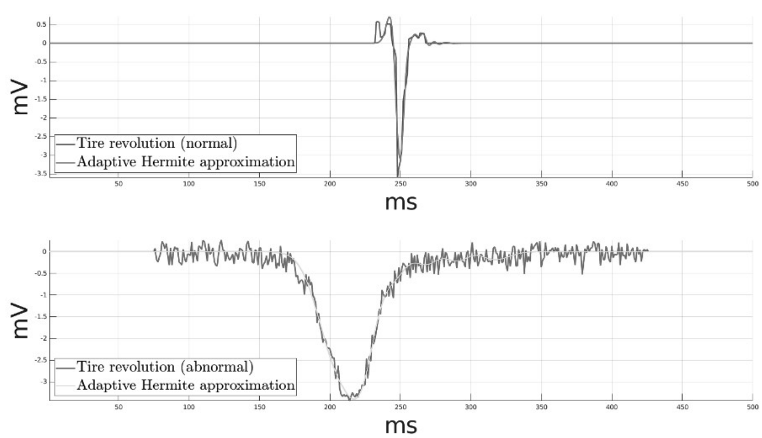

Figure 12. TOP: high speed signal and adaptive Hermite approximation. BOTTOM: low speed signal and adaptive Hermite approximation. Adaptive Hermite functions can model shifts along the time scale and in support width.

2.2 Variable projection operators

In practice, instead of $f \in L_2(R)$, we only have access to the signal $\mathbf{f} \in R^N$, where N is a natural number. Each component of $\mathbf{f}$ can be considered a sampling of f at some point $x \in R$. In order to model these discrete signals, consider the matrix $\Phi^{\lambda,\tau} \in R^{N\times n}$, whose columns contain an equidistant sampling of the first n adaptive Hermite functions. The output of the adaptive high pass filter described at the beginning of this section can then be written as

$$T^{\lambda,\tau}(\mathbf{f}) := \mathbf{f} - \Phi^{\lambda,\tau}\left(\Phi^{\lambda,\tau}\right)^{+}\mathbf{f}, \tag{5}$$

Where $\left(\Phi^{\lambda,\tau}\right)^{+}$ denotes the pseudo inverse of $\Phi^{\lambda,\tau}$. The transformed signals (5) can also be interpreted as the part of the discrete signal $\mathbf{f}$ which falls onto orthogonal complement of the subspace spanned by the first n adaptive Hermite functions. For this reason, the operator $P_{\Phi^{\lambda,\tau}}(\mathbf{f}) := \Phi^{\lambda,\tau}\left(\Phi^{\lambda,\tau}\right)^{+}\mathbf{f}$ is usually referred to as a variable projection operator [14, 15]. We note that identifying the optimal nonlinear parameters λ and τ for the transformation (5) leads to the optimization problem

$$\min_{\lambda>0,\tau\in R} r_2(\lambda,\tau) = \min_{\lambda>0,\tau\in R} \|\, \mathbf{f} - \Phi^{\lambda,\tau}\left(\Phi^{\lambda,\tau}\right)^{+}\mathbf{f}\|_2^2 \tag{6}$$

Based on the work of Golub and Pereyra [14], the partial derivatives of the functional r_2 in the above formula can be explicitly given and gradient based methods can be applied to solve (6).

2.3 Variable projection networks

Since the partial derivatives with respect to the dilation and translation parameters can be calculated in (5), it is feasible to implement (5) as a layer in a neural network. This idea was first introduced in [10]. In general, neural networks, whose first layers implement adaptive orthogonal transformations such as the ones described in (5) are referred to as variable projection neural networks (VP-NET) [10]. In VP-NET, the parameters of these adaptive transformation (or variable projection) layers are optimized together with the weights of an underlying neural network. The variable projection layers are responsible for learning an appropriate representation of the data, while the underlying neural network learns to classify the data examples. Another popular approach for adaptive data representation is the use of convolution layers. Although there are important similarities between the two architectures [10], unlike convolution layers the learned parameters in a VP-NET are interpretable to humans, provided the orthogonal transformations implemented by the variable projection layers were chosen carefully. In our case, using variable projection layers implementing transformation (5), the learned parameters are highly correlated to the speed of the vehicle. For a more thorough review of the VP-NET architecture, we refer to [10].

3. Embedded hardware implementation on a microcontroller

In addition to the interpretability of the learned parameters, variable projection layers provide additional benefits. It has been shown that convolution layers producing effective data representations contain more free parameters [1, 10]. This makes VP-NET based neural networks more appropriate for use with low-capacity hardware. In a real-world application, the VP-NET classifier would have to run on low-cost and low-capacity hardware such as a microcontroller or field programmable gate array (FPGA). This is especially true, if we would like to propose a measurement system that is feasible for commercial use, where production and replacement costs must be considered. Since neural network based machine learning methods have proven to be useful in solving environmental recognition problems [1, 5], embedded implementation of such algorithms has enjoyed significant attention in recent years.

Introduction of AI libraries optimized for low-capacity hardware such as Tensorflow Lite [16] and STM Cube-AI [17] have been useful in the creation of embedded ML applications.

For the first time in this paper, we demonstrate that the VP-NET realization proposed in [1] can be implemented on an STM32 F401RE microcontroller. This is an important step towards the real time application of our measurement system. Figure 4. illustrates the hardware used to implement our proof-of-concept application, which solves the above-described road abnormality recognition problem. The STM32 F401RE microcontroller fulfills the above-mentioned criteria of cost effectiveness, as at the time of writing this manuscript its price is around 13 EUR. Taking into account the hardware specifications of the read-out electronics, force sensors and the proposed implantation methodology [1], we can assume that vehicle or tire manufacturers could incorporate our system into their products in a cost-effective manner. The most difficult challenge associated with implementing the VP-NET classifier on the above-mentioned microcontroller is overcoming the memory and computational limitations. Indeed, the comes with an ARM 32 STM32 F401RE -bit Cortex-M4 CPU with a maximum frequency of 84 MHz and 512 kB of flash memory. These resources are significantly less than the usual environment in which neural network based classifier models are usually used. Below, we describe the specifics of our VP-NET implementation, including our pre-processing steps and architecture choices.

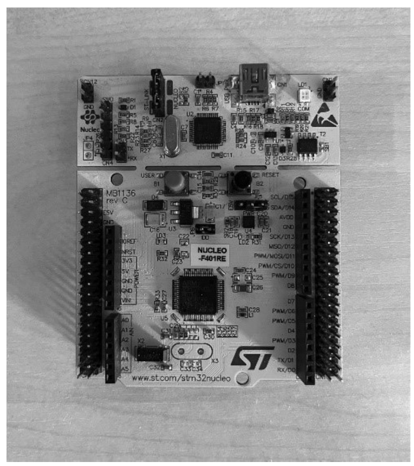

Figure 13. The STM32 F401RE microcontroller used to implement our proposed model-driven classifier.

3.1 Model architecture

In order to implement the road abnormality detection application using VP-NET on a microcontroller, the following steps need to be carried out. First, a VP-NET classifier, similar to the one proposed in [1] has to be trained on a previously gathered dataset. This step of the application has to be done offline for the following reasons. Firstly, training a neural network (even as small as the VP-NET proposed in [1]) consumes considerable amounts of resources. Secondly, construction of the training set can significantly influence the generalization power of the proposed classification model. For the road surface abnormality application for example, the following problem has to be considered. In reality, we can assume that most encountered surfaces will not contain major abnormalities such as large potholes. This means, that if the classification model was trained in real time, most of the input data segments influencing the parameters of the neural network would correspond to normal road surfaces. Thus, inevitably, the classifier would become biased against road abnormalities. For the above reasons, a proper classification model capable of recognizing abnormal road surfaces has to be trained on a balanced data set, which contains roughly the same amount of normal and abnormal data examples. In this study, we used the dataset described in [1] to train our VP-NET model offline. This dataset consisted of 282 normal and 235 abnormal tire revolutions.

The VP-NET model proposed in [1] consisted of a single variable projection layer which implemented transformation (5) using $n = 11$ adaptive Hermite functions. This was followed by three fully connected layers [18] and a final layer containing a single neuron and sigmoid activation. ReLu activation layers were placed between the variable projection layer and the first fully connected layer as well as between each subsequent fully connected layer. It is worth noting, that except for the ReLu and sigmoid activation functions, each layer could be implemented using matrix-vector arithmetic. Indeed, the output of the variable projection layer is given by the matrix operations in equation (5), while each fully connected layer implements the linear transformation

$$\mathrm{W}\mathbf{x} + \mathbf{b}, \quad \left(\mathrm{W} \in \mathrm{R}^{\mathrm{m}\times\mathrm{M}}, \mathbf{x} \in \mathrm{R}^{\mathrm{M}}, \mathbf{b} \in \mathrm{R}^{\mathrm{m}}\right) \tag{7}$$

where $\mathbf{x}$ denotes the input of the layer. As mentioned above, the learnable parameters λ, τ of the variable projection layer and the weights and biases W and $\mathbf{b}$ of the fully connected layers were set during an offline training phase. Each fully connected layer of the state-of-the-art VP-NET proposed in [1] consisted of 64 neurons. This means, that the size of the matrix W was $\mathrm{N} \times 64$ for each layer, where N denoted the dimension of the input tire revolution signals $\mathbf{f}$. In the dataset used to train the models proposed in [1], each tire revolution was zero padded to match the length of 500 data points, or in other words $\mathbf{f} \in \mathrm{R}^{500}$. We already mentioned in section 2, that the effective support of the tire revolution signals changed with the vehicle's speed. The length of 500 data points was chosen to ensure that every measured data point of tire revolutions corresponding to very low speeds could be represented without down sampling. In our current application, however, we could greatly reduce the size of the matrices involved in evaluating a single tire revolution by reducing the size of the vectors representing the inputs. Since less than third of the total number of tire revolutions exceeded the size of 128, this size reduction could be done for most signals by removing the artificially added zero padding. For the few remaining signals, whose support exceeded the length limitation, we applied a down sampling operation. The new dataset created in this way contained a total of 96 down sampled signals. The benefit of reducing the length of the data examples was that the forward pass operation of each layer of the VP-NET could now be described using smaller matrices. In addition to reducing the size of the input examples, we also reduced the number of neurons in each fully connected layer of the VP-NET proposed in [1]. We sequentially reduced the number of neurons in each layer, while the mean accuracy of the trained light weight VP-NET (on the test set) stayed above 96% using 5-fold cross validation. This threshold was chosen, because this was the accuracy of the second best performing convolutional model in [1] considering the same, 5-fold cross validation testing. When determining this threshold we used the modified data set containing reduced length signals. Using this technique, we settled on a simplified light weight VP-NET architecture with a single variable projection layer followed by three fully connected layers each containing 20 neurons. The total number of elements appearing in the matrices needed to implement the forward pass for each layer is illustrated in Table 1.

Table 1. Total number of elements in the matrices implementing each layer

	Proposed light weight VP-NET	VP-NET (state-of-the-art as appearing in [1])
Variable projection layer	2816	11000
1st fully connected layer	2560	32000
2nd fully connected layer	400	4096
3rd fully connected layer	400	4096
Total number of elements	6176	51192

Even though the above-described simplification of the model slightly reduces performance, the accuracy with which the proposed lightweight VP-NET recognizes road abnormalities remains convincing. We further discuss our method of selecting the hyper parameters in section 4. In our numerical experiments, we show that this model and data size reduction results in acceptable computational speed and our implementation can be used to conduct real time experiments in the future. We note that the state-of-the-art architecture proposed in [1] could also be used in embedded hardware implementations, however in this case we advise to use a more computationally capable (and therefore more expensive) microcontroller than our test hardware. A schematic representation of the proposed reduced VP-NET classifier is illustrated in figure 5.

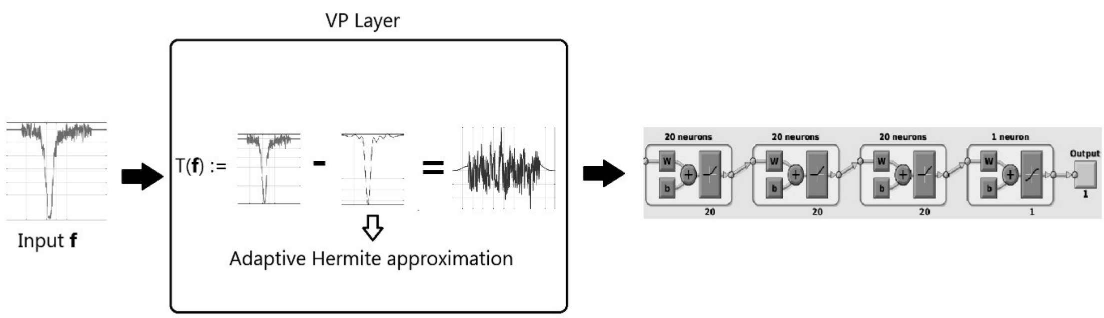

Figure 14. Schematic of the proposed model-driven VP-NET classifier. The first layer implements the adaptive orthogonal transformation (5).

3.2 Notes on the implementation

Since the implemented model-driven network architecture could be described fully with matrix-vector multiplications, additions and element wise applications of the activation functions, we decided to use the Eigen [19] C++ library as the basis of our implementation. Even though libraries specializing at implementing previously trained models on low capacity hardware exist (see e.g. [16] or [17]), these usually work by converting a model using a popular deep learning python library into efficient C code. In most cases however, these libraries have difficulty converting unconventional neural network layers such as the variable projection layer proposed in our application. This, and the simple architecture of the proposed network prompted us to rely on our own, native implementation instead. In our current experimental setup, the evaluation of a single tire revolution consists of the following steps:

- A tire revolution signal (see figure 1.) is serialized using STM's HAL library [20] and transmitted to the microcontroller via USB cable
- The microcontroller evaluates the data example by conducting the steps depicted in figure 5.
- The microcontroller responds with 1 (in case the data example was identified as an abnormal revolution) or 0 (in case the input was deemed normal) via HAL protocol.

4. Numerical experiments

Our experiments were designed to satisfy the two important criteria regarding our microcontroller-based implementation of the proposed reduced variable projection network. Firstly, we wanted to demonstrate that the proposed classification scheme (see section 3) provided similar accuracy on our dataset to the state-of-the-art. More precisely, our objective was to reduce the number of weights associated with the fully connected layers of the VP-NET, while ensuring that model precision remained close to that of the state-of-the-art. In fact, we aimed to propose a simplified variation of the VP-NET which still outperformed the rest of the classifiers investigated in [1]. The second best performing traditional neural network architecture in [1] was a simple convolution based architecture (CNN). This network contained a convolution layer with a single filter using a kernel size of 25. To make the results of our experiments comparable to the results detailed in [1], we had to contend with the fact that we modified the original dataset proposed in [1]. Indeed, as explained in section 3, we down sampled the tire revolution signals to match the maximum length of 128 in order to further reduce the size of the matrices needed to evaluate a single data example. In order to make the comparison fair, we re-ran the 5-fold cross validation experiment on the CNN architecture on the modified dataset as well. All models in our experiment were trained on the new, modified datasets (with zero padding removed to and down sampling added to ensure a length of 128 samples for each tire revolution). Our results are demonstrated in Table 2.

Table 2. Results of 5-fold cross validation on the original and proposed datasets

Model name	**Mean accuracy**	**Highest accuracy**	**Lowest accuracy**
CNN	95.93%	98.08%	94.17%
VP-NET (as proposed in [1])	97.09%	100.0%	94.17%
Proposed light weight VP-NET	96.12%	98.05%	94.23%

Table 2 shows that in terms of accuracy, the proposed light weight VP-NET architecture is only outperformed in terms of accuracy by the state-of-the-art (VP-NET with larger fully connected layers than the one proposed in this paper). On the other hand, the proposed architecture outperforms the second most capable CNN. In addition to the improvement in accuracy, we would like to note that implementing the convolution layer as a matrix operation can be formalized as

$$\mathbf{y} = \mathrm{Q}\mathbf{f}, \qquad \left(\mathrm{Q} \in \mathrm{R}^{\mathrm{m}\times\mathrm{N}}, \mathbf{y} \in \mathrm{R}^{\mathrm{m}}\right) \tag{8}$$

where Q denotes a Toeplitz matrix of size $\mathrm{m} \times \mathrm{N}$ and m is the length of the discrete convolution filter while N denotes the dimension of the tire revolution signals. In other words, considering the simplified dataset proposed in section 3, our light weight VP-NET architecture can perform the Hermite function based adaptive orthogonal transformation (5) using two matrices of size 128×11, while the CNN architecture needed a matrix operation of size 25×128. Finally, we note, that the CNN and state-of-the-art VP-NET models show slightly lower accuracies than the results we achieved in [1]. This is due to the effect of our down sampling transformation applied to "slow" tire revolutions. We conclude that the proposed VP-Net architecture is not only more accurate, but also more computationally efficient than the second best performing convolution network.

The second important criteria that our implementation had to meet was a quick enough response time. To verify this, we measured the time needed for the microcontroller to provide us a response on every tire revolution (517 signals in total). Our results can be seen in Table 3.

Table 3. Response times of the microcontroller implementing our light weight VP-NET classifier.

Average response time (s)	**Shortest response time (s)**	**Longest response time (s)**
0.047	0.015	0.156

From table 3, we can deduce, that the average response time was about 0.047 seconds, with negligible deviation. Since our test vehicle was equipped with size 205/55R16 tires with a circumference of around 2 meters, this result means that the vehicle could in theory recognize an encountered road abnormality around 0.047 seconds after passing it. Considering a vehicle speed of 50 km/h, this means that the road abnormality would fall into a ~ 0.65 meter radius of the vehicle's position at the time of reporting. This proximity is low enough to map "problematic" parts of the road network with high accuracy, if the information gathered this way is collected in a database [3]. This can contribute to reducing maintenance costs. In fact, these results show that at the velocity of 50 km/h any road abnormality detected during a tire revolution would get reported before the next tire revolution is finished. Furthermore, we highlight that in these experiments we intentionally chose the STM32 F401RE low cost and low-capacity microcontroller. In a real-world application, it is feasible to think that the proposed methods would be implemented on a higher capacity device. Our results thus demonstrate that the proposed intelligent tire sensor-based measurement system could be used in real-world applications such as the recognition of road surface abnormalities.

5. Current research directions and challenges

Several methods already exist for detecting road surface abnormalities. These approaches rely on signals from accelerometers attached to, or placed inside the chassis [3], or images acquired from a camera mounted onto the vehicle [4, 5]. Our intelligent tire sensor system addresses several issues associated with these approaches. In contrast to the accelerometer-based approaches, where the vibrations recorded by the sensor as often damped by car suspension and other factors, our sensors record data as close to the road surface as possible. In addition, separate sensors can be implanted into each tire, therefore we can describe the road surface with finer resolution. When compared to image processing-based methods, the main advantage of the proposed sensor system is that changes in the road surface composition can be detected even with limited visibility (such as at night or in bad weather conditions).

Our current research target is to create a proof-of-concept application capable of recognizing road surface abnormalities in real time. The most important step towards this goal is the main result of this paper, namely the implementation of our data-driven signal processing algorithms on low-capacity hardware. In order to create a product for commercial use, however, several issues still need to be addressed. Firstly, as detailed in the previous sections, the signal processing part of our measurement system expects signals measured during full rotations of the tire. This means, that any real-life application should segment the measurements appearing on the CAN network of the vehicle into full tire rotations. Ideally, the segmentation would be done using sensors capable of measuring the wheel angles, or at least capable of indicating that a full rotation of a tire has taken place. We are in the process of equipping our test vehicle with such capabilities. Another interesting research objective is the construction of light weight classification algorithms employing adaptive orthogonal transformations, without the expensive evaluation process of the fully connected neural layers in VP-NET. One candidate algorithm of interest is the implementation of VP-SVM [21] on microcontrollers which extends the usual support vector machine (SVM) classifier objectives with adaptive orthogonal transformations.

Since our sensors measure the forces caused by tire wall deformation, it is also reasonable to investigate how our measurements correlate to specific maneuvers undertaken by the vehicle. In particular, we are interested in estimating certain parameters describing the dynamical state of the vehicle based on the tire sensor signals. Among these parameters, of central interest would be the ones describing sudden changes in vehicle dynamics (e.g. emergency braking, loss of traction etc.) such as the road tire friction coefficient [6]. Because the relationship between any such parameters and the tire sensor output signals is highly nonlinear and difficult to model directly, we continue relying on the hybrid, (model and data-driven) methods such as VP-NET. This of course requires the accumulation of reference data therefore we are currently in the process of specifying and conducting measurements related to these objectives.

Conclusion

In this paper, we further developed our intelligent tire based measurement system introduced in [1]. Specifically, we proved through an example application, that the proposed model-driven neural network architecture can be implemented on a microcontroller to efficiently identify abnormalities on the road surface. We conducted numerical experiments to justify our implementation choices and showed that the proposed light weight version of the state-of-the-art classifier [1] can be used to solve this problem in real time. We compared our measurement system to other popular approaches with respect to road abnormality detection and concluded, that the use of our proposed method has several benefits (finer resolution of the road surface, usable regardless of visibility conditions). We discussed the current direction and challenges of our research and outlined our future plans.

Acknowledgement: *We would like to thank Attila Nagy for contributing to this work with sensor integration. Project no.\ C1748701 has been implemented with the support provided by the Ministry of Culture and Innovation of Hungary from the National Research, Development and Innovation Fund,* *financed under the NVKDP-2021 funding scheme. This project was supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences. The research was supported by the European Union within the framework of the National Laboratory for Autonomous Systems. (RRF-2.3.1-21-2022-00002). This research was funded by the Ministry of Innovation and Technology of Hungary from the National Research, Development and Innovation Fund, financed under the TKP2021 funding scheme, grant number TKP2021-NVA-03*

References

[1] T. Dózsa, J. Radó, J. Volk, Á. Kisari, A. Soumelidis and P. Kovács, "Road Abnormality Detection Using Piezoresistive Force Sensors and Adaptive Signal Models," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1-11, (2022)

[2] Xiong, Yi, and Ari Tuononen. "Rolling deformation of truck tires: Measurement and analysis using a tire sensing approach." *Journal of Terramechanics* 61 (2015): 33-42.

[3] Du, Ronghua, et al. "Abnormal road surface recognition based on smartphone acceleration sensor." *Sensors* 20.2 (2020): 451.

[4] Vigneshwar, K., and B. Hema Kumar. "Detection and counting of pothole using image processing techniques." *2016 IEEE International Conference on Computational Intelligence and Computing Research (ICCIC)*. IEEE, 2016.

[5] Pandey, Anup Kumar, et al. "Convolution neural networks for pothole detection of critical road infrastructure." *Computers and Electrical Engineering* 99 (2022): 107725.

[6] Erdogan, Gurkan, Lee Alexander, and Rajesh Rajamani. "Estimation of tire-road friction coefficient using a novel wireless piezoelectric tire sensor." *IEEE Sensors Journal* 11.2 (2010): 267-279.

[7] Zhu, Bing, Jiayi Han, and Jian Zhao. "Tire-pressure identification using intelligent tire with three-axis accelerometer." *Sensors* 19.11 (2019): 2560.

[8] Basargan, Hakan, et al. "Road quality information based adaptive semi-active suspension control." *Periodica Polytechnica Transportation Engineering* 49.3 (2021): 210-217.

[9] J. Radó, A. Nagy, and J. Volk, "Measuring device for kinetic dynamic state of a vehicle tyre, gumiabroncs menetdinamikai állapotát eszköz," Patent 1 900 189, Oct. 3, (2019).

[10] Kovács, Péter, Gergo Bognár, Christian Huber, and Mario Huemer. "VPNET: Variable Projection Networks." *International Journal of Neural Systems* (2021).

[11] Dózsa, Tamás, and Péter Kovács, "ECG signal compression using adaptive Hermite functions." In *ICT Innovations 2015: Emerging Technologies for Better Living 7*, pp. 245-254., (2016).

[12] Gautschi, Walter. Orthogonal polynomials: computation and approximation, OUP Oxford, 2004.

[13] Szegő, Gábor. Orthogonal polynomials. Vol. 23. American Mathematical Soc., 1939.

[14] Golub, Gene H., and Victor Pereyra. "The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate." *SIAM Journal on numerical analysis* 10.2 (1973)

[15] Golub, Gene, and Victor Pereyra. "Separable nonlinear least squares: the variable projection method and its applications." *Inverse problems* 19, no. 2 (2003)

[16] Martín Abadi, Ashish Agarwal, et. al. „TensorFlow: Large-scale machine learning on heterogeneous systems", (2015). Software available from tensorflow.org.

[17] Falbo, Vincenzo, Tommaso Apicella, Daniele Aurioso, Luisa Danese, Francesco Bellotti, Riccardo Berta, and Alessandro De Gloria. "Analyzing machine learning on mainstream microcontrollers." In *Applications in Electronics Pervading Industry, Environment and Society: APPLEPIES 2019 7*, pp. 103-108. (2020).

[18] J. Schmidhuber, "Deep learning in neural networks: An overview", *Neural networks* 61 (2015)

[19] Gael Guennebaud, Benoit Jacob et. al. „Eigen v3", *http://eigen.tuxfamily.org,* (2010).

[20] STMicroelectronics, "STM32Cube MCU Package for STM32F4 series," *https://www.stm.com*, (2020).

[21] T. Dózsa and P. Kovács, „Variable projection support vector machines", *Proc. 4th International Conference on Advances in Signal Processing and Artificial Intelligence (ASPAI)*, Corfu, Greece, Oct. 19–21, 2022, pp. 47–52. (2022)