

Deep Learning Towards Intelligent Vehicle Fault Diagnosis

Mohammed Al-Zeyadi*, Javier Andreu-Perez*, Hani Hagras*, Chris Royce†, Darren Smith†, Piotr Rzonsowski†, Ali Malik‡

*School of Computer Science and Electronic Engineering, University of Essex, United Kingdom

{ma18964, javier.andreu, hani}@essex.ac.uk

†Cognitran Limited, United Kingdom

{chris.royce, darren.smith, piotr.rzonsowski}@cognitran.com

‡School of Electrical and Electronic Engineering, Technological University Dublin, Ireland
ali.malik@tudublin.ie

Abstract—Recently, the rapid development of automotive industries has given rise to large multidimensional datasets both in the production sites and after-sale services. Fault diagnostic systems are one of the services that the automotive industries provide. As a consequence of the rapid development of cars features, traditional rule-based diagnostic systems became very limited. Therefore, more sophisticated AI approaches need to be investigated towards more efficient solutions. In this paper, we focus on utilising deep learning so as to build a diagnostic system that is able to estimate the required services in an efficient and effective way. We propose a new model, called *Deep Symptoms-Based Model* Deep-SBM, as an approach to predict a wide range of faults by relying on the deep learning technique. The new proposed model is validated through a set of experiments in order to demonstrate how the underlying model runs and its impact on improving the overall performance metrics. We have applied the Deep-SBM on a real historical diagnostic data provided by Cognitran Ltd. The performance of the Deep-SBM was compared against the state-of-the-art approaches and better result has been reported in terms of accuracy, precision, recall, and F-Score. Based on the obtained results, some further directions are suggested in this context. The final goal is having fault prediction data collected online relying on IoT.

Index Terms—AI, deep learning, deep neural network, vehicle fault diagnosis, Internet of Things (IoT).

I. INTRODUCTION

In today's intelligent systems, AI technologies have been widely applied in the automotive industry, both in the production sites and after-sale services. The fault diagnostic system is one of the services the automotive industry provides. Any improvement in the fault diagnostic system will facilitate faster repairs which lead to cost savings and improved customer satisfaction thereby, improving the quality of the overall product and services which in turn will increase brand loyalty. Technically, faults diagnostic system consists of two stages [1]. First, *On-Board Diagnostics*, where the Electrical Control Units (ECUs) identify issues with the vehicle's performance and register pre-identified code that can be used for the later diagnosis. Second, *Off-Board Diagnostics*, where a Repair and Maintenance Information (RMI) system is used for the purpose of fault detection after collecting the data from the first stage. The second stage is longer as it requires significant manual effort and domain knowledge.

In automotive vehicles, ECU is an embedded system that controls one or more of the electrical systems/subsystems in a vehicle. It works as an early stage self-diagnosis to detect and record abnormal behaviours of vehicle components. The number of ECUs are directly proportional to the complexity, i.e. the greater the number of ECUs, the more complex the diagnosis. Modern luxury vehicles are equipped with a large number of ECUs, for instance BMW 7-series has 150 automotive ECUs¹. As a consequence, the traditional hand-crafted rule-based diagnostic systems become inaccurate and time-consuming, as will be described in Section V. Therefore, to address this challenge, a smart diagnostic system with the help of AI technologies is required and needs more investigation.

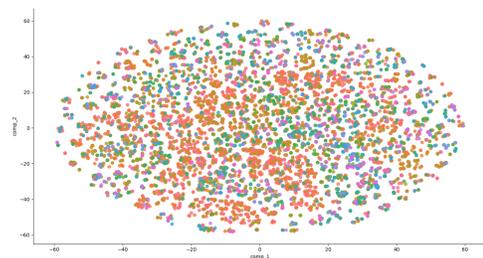


Fig. 1: A t-SNE projection for 100k samples from diagnostic session data in 2018, dots represent samples from different fault types.

This paper deals with the problem of fault diagnostics where the scope of fault types is up to 3000. This large number of fault types makes the traditional diagnostic systems impractical as the extraction of the requisite discriminative features are very complex as shown in Figure 1, which has been conducted using t-distributed stochastic neighbor embedding (t-SNE) [2]. In this regard, deep learning is a promising solution to the aforementioned problem to overcome the challenges of the traditional machine learning techniques [3].

In fact, deep learning is an important building block for learning the requisite discriminative features directly from a raw data. Traditionally, vehicle's diagnostic systems were

¹<https://autotechinsight.ihsmarket.com/>

highly dependant on human knowledge and handcrafted expert rules. However, this is not feasible with the current complex vehicle systems anymore. Thereon, deep learning is one of the possible solutions that could automate the diagnostic process. Deep learning has also been proposed as a solution to resolve various kinds of vehicle-related problems such as self-driving, speed prediction, vehicle-sensing and tracking [4]–[7].

The rest of the paper is organised as follows. Section II introduces various fault diagnostic techniques from the literature. Section III introduces the problem statement and the contributions of this paper. We then present our learning architecture model and framework in Section IV. Section V presents the experiments and results. Finally, the conclusions and future work are provided in Section VI.

II. RELATED WORK

Car systems are plagued by a number of faults. Since the beginning of the 20th century, the number of faults have sharply increased due to the massive features of the current car system technologies such as GPS, self-driving, bio-metric vehicle access and so on. Vehicle fault diagnosis systems are one of the most essential requirements that increases the reliability of the vehicle through monitoring the state of electronic control systems. A considerable amount of research has been conducted to consider the challenge of fault diagnostic of vehicles in different perspectives. In this section, we discuss some recent studies which are related to the proposed approach of this work.

Recently, AI techniques have become increasingly applied in vehicles fault diagnosis due to the significant impact [8]. In this regard, the authors in [9] proposed a self-diagnosis system for autonomous vehicles that aims to improve the self-diagnosis speed and reduce the overhead. The proposed system consists of three modules where the first module is responsible for data gathering from autonomous vehicles using the Internet of Things, IoT, the second module is an optimised deep learning to initiate a training dataset based on the collected data of first module and finally, the third module is an edge computing based self-diagnosis service. Also, the authors presented the Lightweight In-Vehicle Edge Gateway (LI-VEG) [10] for the self-diagnosis system, which provides rapid and accurate communication between vehicle and self-diagnosis module.

The authors in [11] introduced a new fault detection and diagnosis approach by combining two techniques, these are: the Auto-Associative Neural Networks (AANN) and the Adaptive Neuro-Fuzzy Inference Systems (ANFIS). The proposed approach was tested on a 10-variables vehicle model only. In addition, the same research team presented a new approach in [12] to detect vehicle’s faults in an online fashion based on the historical reported data. The proposed method consists of two phases, an auto-associative neural network and a fuzzy system. The study has demonstrated the ability of the proposed method to detect the faulty variables along with its time of occurrence. Furthermore, the same authors have presented a new framework for online vehicle diagnosis and fault detection

[13]. The presented framework consists of two phases: first phase is an auto-associative neural network that infers the residuals between the incoming data from outsourcing and the learned behaviour of normal operation; second phase is a multi-class Support Vector Machine (SVM) that classifies the reported faults type and the occurrence time. The experiments were conducted on ten variables vehicle monitoring.

The Generative Adversarial Network (GAN) was introduced in [14] as a technique to prevent the risks of vehicles component failures. The study showed that a more robust predictive mechanism can be obtained in order to be used for predictive maintenance.

The authors in [15] proposed a mathematical model based deep learning for the problem of on-board fault diagnosis in high-speed railways. The study showed that the proposed method outperformed both the k-nearest neighbor (KNN) and the artificial neural network with back propagation (ANN-BP) in terms of fault diagnosis.

The authors in [16] presented a fault detection and diagnosis method by using the artificial neural network (ANN) based AC–DC converter. The study showed that different fault properties like severity can be detected.

The authors in [17] presented a new acoustic based fault diagnosis technique of the three-phase induction motor. The study demonstrated the use of the nearest neighbour and the modified classifier based on words coding for classification and recognition. The conducted experimental results showed that the acoustic based fault diagnosis technique has low operation cost and high dependability.

Although these studies provide important insights into the area of fault diagnostics assessment, such studies remain narrow in focus dealing only with certain type of faults at a time without taking into account multiple types of fault. Therefore, in this paper, we focus on defining a comprehensive fault diagnostic system model that deals with a wide range of fault types and different vehicles. The powerful of such model is to discover the hidden patterns and construct a relation between the features of the cars and symptoms in the historical data. This was not very clear with respect to the based diagnostic systems of the existing works.

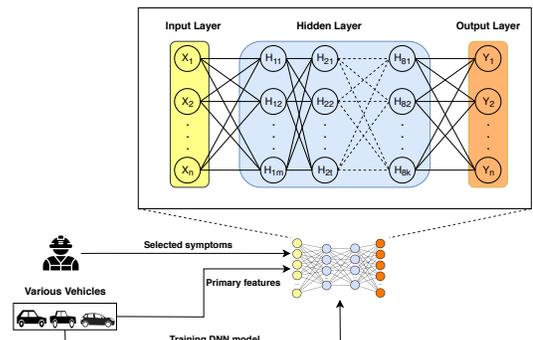


Fig. 2: The architecture of DNN model

III. PROBLEM STATEMENT AND SUMMARY OF CONTRIBUTIONS

We deal with fault diagnostics as a multi-classification problem where each class represents a unique fault type. For this purpose, Deep Neural Network (DNN) is chosen to develop a diagnostic classifier based on historical sessions data. Formally, we define D to be a set of N historical sessions in the form of $\{(x_1, y_1), \dots, (x_N, y_N)\}$, where x_i represents a vector of vehicle features and symptoms, while, y_i represents the fault type. In our case, we dealt with up to 3000 different fault types in the domain of Y . We have the function: $f : X \rightarrow Y$, is learned by using the DNN. The inputs, x_i , to the DNN diagnostic classifier are extracted from the historical diagnostic data in which fault types were identified successfully. A schematic example is provided in Figure 2 as an illustration of classification process.

With the above context in mind, we can summarise the main contributions of this paper as follows:

- **Data parsing:** To generate training and testing data after parsing the historical diagnostic data. This is to set the stage for the creation of DNN diagnostic model.
- **Diagnostic model:** To develop a DNN diagnostic model, which we call Deep Symptom-Based Model (Deep-SBM), that can predict fault types.
- **Experiments:** To examine the proposed model, the performance of the Deep-SBM is compared with another three baseline models. In this context, we provide a set of experiments to test the new proposed model. The conducted results prove that the proposed model improve the performance of fault diagnostic problem.
- **Diagnostic explainer:** To provide a comprehensive understanding upon the diagnostic fault type, we used Local Interpretable Model-agnostic Explanations (LIME) in our proposed model.

IV. THE PROPOSED FRAMEWORK

In this section we discuss the proposed framework and its components. From a high level point of view, Figure 3 illustrates the main components of our proposed framework where the grey coloured component represents the main contribution of this paper.

We discuss the components of this framework in detail.

A. Diagnostic process and session generator

This component is concerned with the different stages of the diagnostic process. In fact, the process of this component is beyond the scope of this paper. However, we have obtained historical session data that represents an audit trail from Cognitran² Ltd. The diagnostic process starts when a vehicle with an unknown malfunction is connected into an RMI for faults detection purposes. RMI can be seen as a guide approach that aids technicians to diagnose a vehicle's issues. Usually, this process starts by selecting the set of symptoms according to the concerns of vehicle owner. Next, the physical features

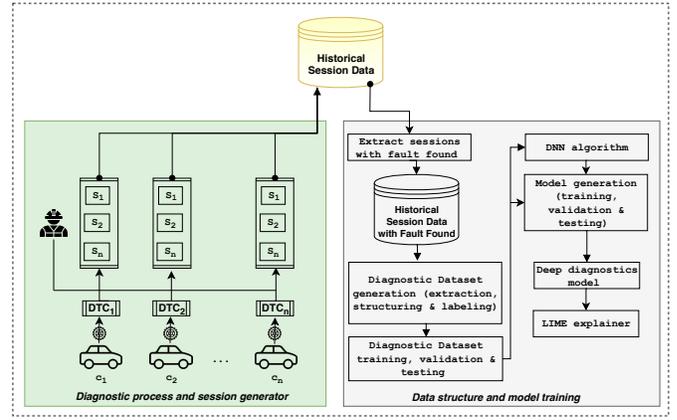


Fig. 3: Proposed Framework

of vehicles are combined with the symptoms. Then, this combination is referenced against an authored *symptom-map* to produce an ordered list of diagnostic tests.

After that, the failure of any diagnostic test will then provide an indicted component, which guides the technician through the repair/post-repair instructions. Finally, all the previous steps are recorded in an audit trail, know as a diagnostic feedback session file. These session files have been parsed to be used as an input for the next component of this framework.

B. Data structure and model training

This component highlights the main contribution of this paper. The first stage is to extract the parsed sessions of the previous component. We only focused on the sessions where the faults have been successfully identified. In other words, we discard every session whose fault was not successfully determined.

Then, the diagnostic dataset is generated by converting the extracted sessions into vectors. We represent the vectors as a set of vehicle features and selected symptoms. Such vectors are used as a input to the proposed Deep-SBM. We faced two challenges. First, these vectors are not linearly separable. Second, as mentioned earlier (Section I), the number of fault types is very large. With these two challenges, it is very clear that the traditional shallow machine learning techniques are inefficient to deal with these problems.

To overcome these challenges, we used a DNN aiming at gaining a better result compared with the traditional machine learning techniques.

C. Deep Learning Model Architecture

The vehicle features and symptoms, which have been represented as vectors, are encoded in a boolean input tensor, where $x_k^i \in \{0, 1\}$, $i \in [1, N]$, $N=2093$. The fault types is given by g_k where G is the total number of fault types and k is one of K sessions. The proposed learning method is based on a deep sequential 9 layer architecture, plus the standard input and output layers. The 9 layer architecture encompasses 3 repetitive blocks. Each block consists of 3 layers as follow:

²<https://www.cognitran.com/>

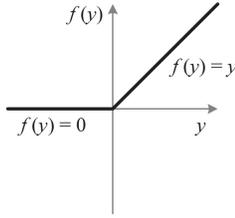


Fig. 4: ReLU function where $f(x)$ is either 0 when $x < 0$, or a linear when $x \geq 0$.

- 1) A regular densely connected neural-network layer with a number of units, Ψ , can be arrived at:

$$\Psi = \delta / (\eta + \beta) * \lambda \quad (1)$$

Where: δ represents the number of tensors, η represents the number of fault types, β represents the length of tensor and $1 \geq \lambda < 5$.

- 2) The activation function for the regular densely connected hidden Layer is a Rectifier Linear Unit (ReLU) [18]. In general, this function can be arrived at:

$$f(x) = \max(0, x), \text{ s.t. } f(x) = \{x \in \mathbb{R} | x \geq 0\} \quad (2)$$

where x is the input to a neuron. This type of activation function works by thresholding values at 0, simply by giving the output of $f(x) = 0$ when the $x < 0$ and conversely, it outputs a linear function when $x \geq 0$. Figure 4 depicts an illustration of the ReLU activation function.

- 3) A Dropout layer with a dropping rate equal to 20% for a good regularisation and over fitting prevention [19]. Figure 5, illustrates an example about multi-layer neural network with and without dropout.

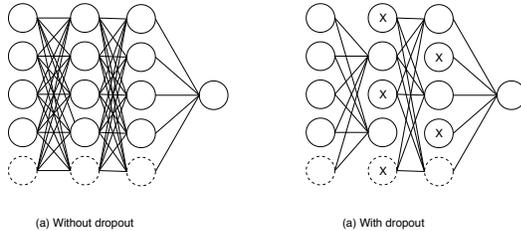


Fig. 5: Dropout example

Finally, the three blocks are connected to a dense Softmax layer for final classification. The model schema architecture is giving in Table I. As an optimiser function for our training model, ADAM method was utilised for stochastic optimisation of DNN parameters, which minimises a categorical cross-entropy function between the training and the predicted diagnostic tactic ($x^k \rightarrow y^k$) where k is a specific session.

D. Understanding Diagnostic Prediction by Estimating Inputs Relevance using LIME

In this paper we have used LIME algorithm [20] to give an interpretable, locally faithful explanations of the individual

TABLE I: The proposed Deep Learning Model schema architecture

Layer (type)	Output Shape	Param #
dense_1 (Dense)	(None, 2093)	4382742
activation_1 (Activation)	(None, 2093)	0
dropout_1 (Dropout)	(None, 2093)	0
dense_2 (Dense)	(None, 480)	1005120
activation_2 (Activation)	(None, 480)	0
dropout_2 (Dropout)	(None, 480)	0
dense_3 (Dense)	(None, 240)	115440
activation_3 (Activation)	(None, 240)	0
dropout_3 (Dropout)	(None, 240)	0
dense_4 (Dense)	(None, 1178)	283898
Total params: 5,787,200		
Trainable params: 5,787,200		
Non-trainable params: 0		

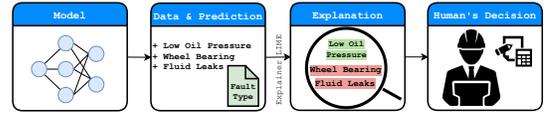


Fig. 6: Explaining individual predictions

diagnostic prediction made by Deep-SBM model. This is to help with answering the main crucial question of any diagnostic ML model, which is *how trustworthy is the prediction*. This component is concerned with answering this question by supporting the diagnostic prediction with an understandable explanation, which provides more credibility to the diagnostic system. For example, providing a list of symptoms that led the model for a prediction, would help the technician decide whether or not to trust the predicted fault type as it is illustrated in Figure 6.

In fact, LIME gives some insights about the importance of every input with respect to the determined fault type, i.e. label, which is locally isolated as a binary model. According to [20], the following equation is used to produce the LIME explanation.

$$\xi(x) = \underset{g \in G}{\operatorname{argmin}} \mathcal{L}(f, g, \pi_x) + \Omega(g) \quad (3)$$

According to equation 3, f and x are the original predictor and features, respectively. In our case, f represents the Deep-SBM and x represents the symptoms and features of vehicle. While, g is the explanation model, which in our case a linear model. Π in equation 3 represents the proximity measure between an instance of z to x to define locality around x . Finally, the model complexity of explanation g is measured by Ω . We refer the interested readers to [20] for more details.

V. EXPERIMENTAL AND RESULTS

The experiments in this section were conducted by considering the session data of 2018 only. As a preliminary investigation, we have evaluated the efficiency and the effectiveness of the current RMI for the historical diagnostic session data of the above mentioned period. The result of this evaluation is shown in Table II.

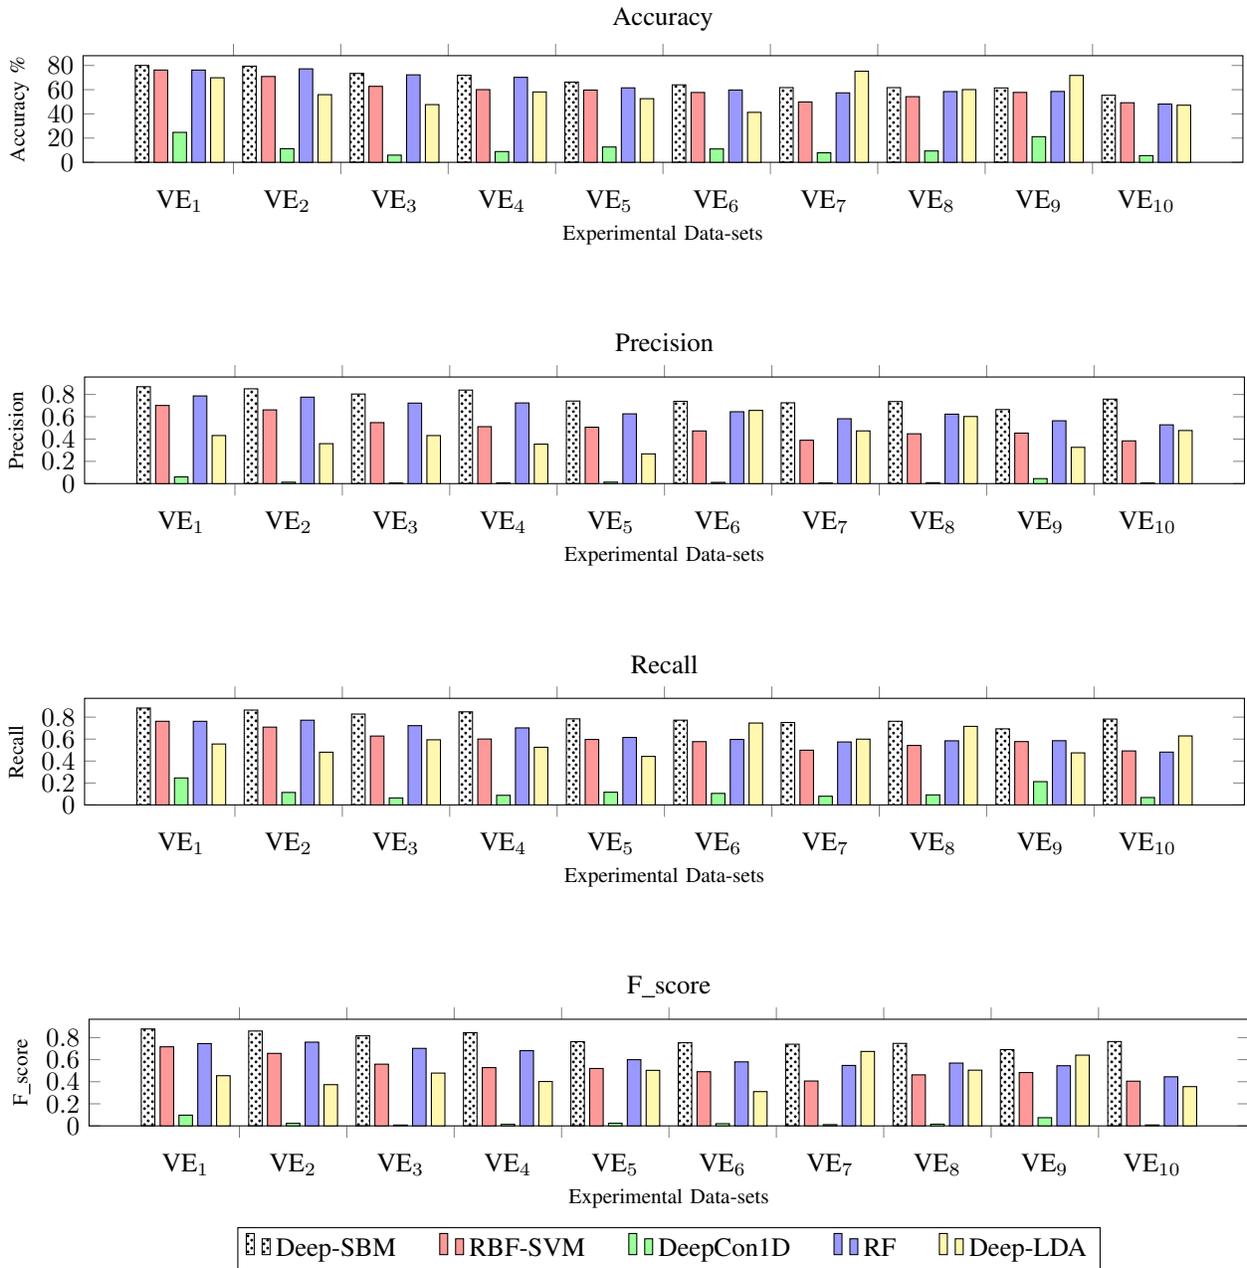


Fig. 7: Performance measurements using 10 different vehicle models

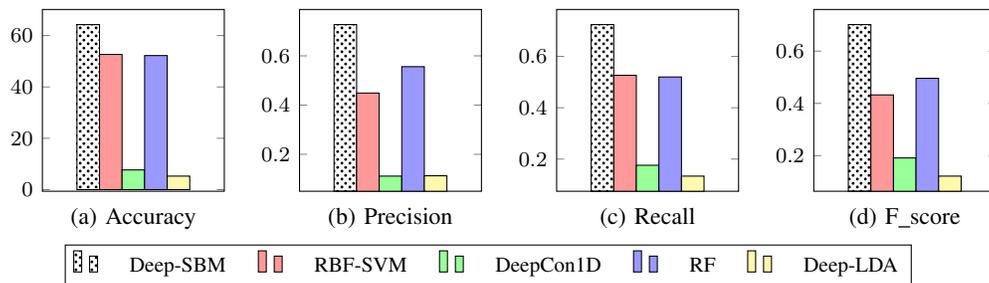
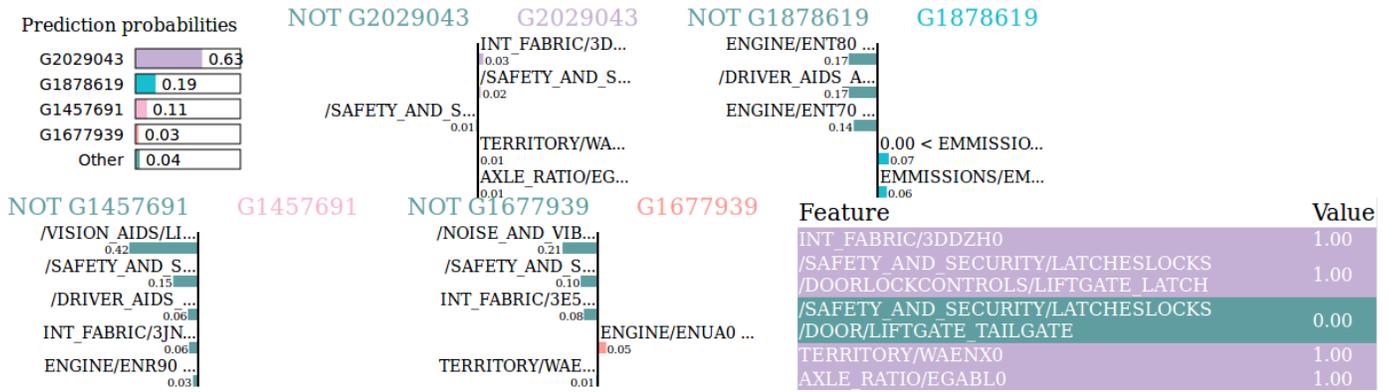
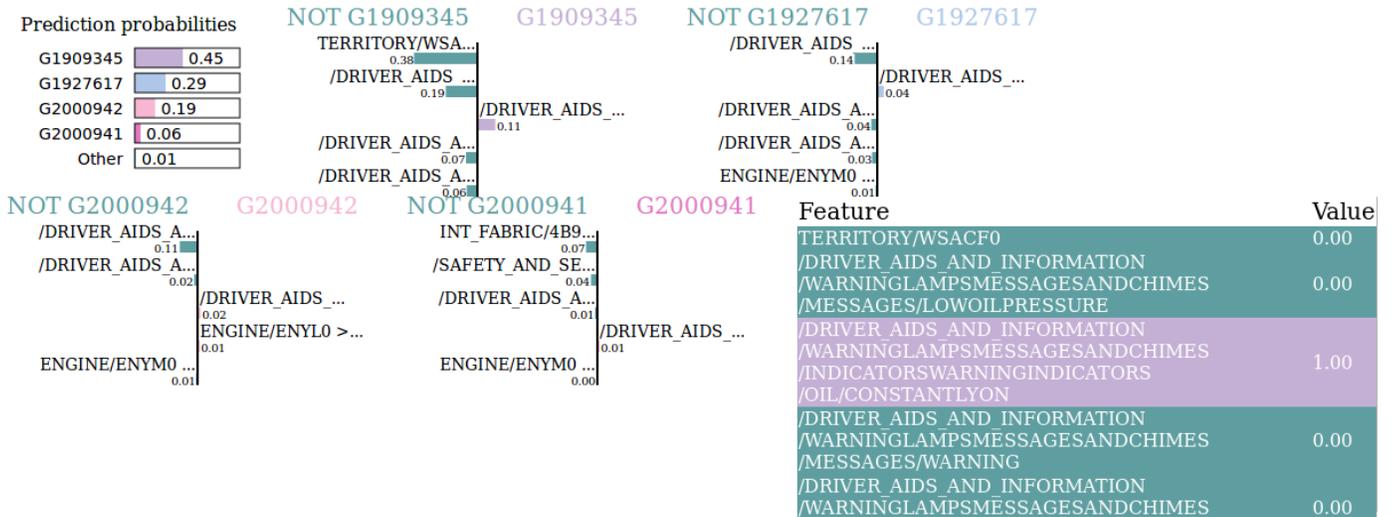


Fig. 8: Performance Measurements using all vehicle models



(a) Example 1



(b) Example 2

Fig. 9: Two examples of LIME Applied to Diagnostic Prediction of Deep-SBM

On one hand, in this evaluation, the accuracy is measured by counting the number of successful and unsuccessful diagnosed sessions. According to Table II, the accuracy on average is less than 50% and this accuracy measures give an impression about the low effectiveness of the current RMI. On the other hand, the time duration is measured by tracking only successfully

diagnosed sessions, which in our case represents the efficiency of the diagnostic system. According to Table II, the reported time duration is too slow in finding fault types. For instance, in January, it takes up to 12 minutes to find fault types, which reflects the inefficiency of the current RMI. This gives another impression about how slow the current system is due to the nature of rule-based diagnostic mechanism.

TABLE II: Statistics of historical diagnosis data in 2018

Month	No. sessions	Accuracy	Avg. diagnostic time
January	121,525	49.42%	12.952
February	116,402	49.30%	11.262
March	116,402	47.75%	11.213
April	99,410	47.16%	10.549
May	108,826	48.83%	10.505
June	103,082	47.66%	10.556
July	115,631	47.73%	11.215
August	116,960	48.65%	12.906
September	115,594	49.07%	12.488
October	127,347	48.14%	12.541
November	99,697	47.41%	12.483
Total	1235542	48.31%	11.743

In terms of evaluation of our proposed model, i.e. Deep-SBM. We compared the Deep-SBM against: 1) Deep Convolution 1D (DeepCon1D), 2) Deep linear analysis discriminant (DeepLDA), 3) Radial Basis Function Support Vector Machine (RBF-SVM), 4) Random Forest (RF). Two experiments were conducted in order to evaluate the performance including four metrics, these are: accuracy, precision, recall and F-score.

Figure 7 shows the first experimental results. This experiment is conducted in order for us to be able to make statements about which model could have the best performance in a small-scale diagnostic domain. To do so, ten vehicle line models were selected; namely $VE_1, VE_2, \dots, VE_{10}$. The total number of the registered fault types is 232, 242, 281, 144, 191, 176,

330, 224, 175, 148, respectively. According to Figure 7, Deep-SBM managed to achieve between [70%-80%] f-scores for all vehicle model lines. In terms of accuracy, Deep-SBM managed to achieve accuracies between [60%-80%] and is the winner in 6 out of 10 vehicle models. In this regard, some of the baseline models were so competitive to the proposed model, for example Deep-LDA. The reason behind that is because we are dealing with the vehicle line models individually whose diagnostic domain range is small-scale. However, the main target of this work is to deal with all vehicle line models in one diagnostic prediction model.

Figure 8, shows the second experimental results. This experiment is conducted with respect to all vehicle line models. In this experimental scenario, we have dealt with 38 vehicle line models and up to 3000 different fault types to be predicted at once. According to Figure 8, the Deep-SBM is far better than the four baseline models with respect to all performance measurements metrics that we used in the first experiment. The architecture of Deep-SBM, which is based on 9 hidden layer plus the standard input and output layers, gives more ability to build a robust diagnostic model that capable to deal with up to 3000 different fault types.

As an additional investigation, we run further a experiment with LIME to assess the relevance of the inputs with respect a determined prediction. Figure 9(a), shows the highest 4 prediction fault types based on its probability, which lies in the top left corner of the figure. Figure 9(a), also shows the features and symptoms that correspond to the highest predictions. Figure 9(b), gives another example of LIME with new symptoms and features. We repeat the experiment of Figure 9(a) with new inputs.

VI. CONCLUSIONS AND FUTURE WORK

In this work a deep learning architecture, Deep-SBM, is proposed to predict a wide range of faults of a vehicle based on the reported symptoms. Deep-SBM is able to accurately classify approximately 3000 different diagnostic fault types. Provided observations consisted on a set of vehicle features and self-reported symptoms, making a total of 2093 inputs. The original experts hand-crafted rule-base (RMI) used by the manufacturer only manages to recognize the specific fault of the problem with a maximum accuracy of 50%.

The performance of the proposed model (Deep-SBM), is tested and evaluated through extensive experiments that have shown how it outperformed the other baseline models. Although some baseline models perform competitively when the domain range is small-scale, such baseline models perform poorly when they are faced with a large-scale domain. However, Deep-SBM still provides more robust and consistent performance in either case, i.e. small and large scale domain with less biased predictions than the other methods within the benchmark. The presented results are very promising for the effort to automatise post-purchase vehicle service and maintenance. It demonstrates a more accurate and robust prediction of the correct fault than state-of-the-art methods used. Nevertheless, the highest impact is still to come when

connected and autonomous vehicles pursue the next step of full autonomy, including self-maintenance and self-service. In future, we will investigate the opportunities that can give us a better understanding of the diagnostic codes so that the predictions of failures are even more autonomous.

REFERENCES

- [1] J. Fish, D. R. Moulton, and K. Gray, "Graphical user interface with on board and off-board resources," Mar. 29 2016, uS Patent 9,299,197.
- [2] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [3] S. L. Oh, J. Vignes, E. J. Ciaccio, R. Yuvaraj, and U. R. Acharya, "Deep convolutional neural network model for automated diagnosis of schizophrenia using eeg signals," *Applied Sciences*, vol. 9, no. 14, p. 2870, 2019.
- [4] M. Daily, S. Medasani, R. Behringer, and M. Trivedi, "Self-driving cars," *Computer*, vol. 50, no. 12, pp. 18–23, 2017.
- [5] Z. Cheng, M.-Y. Chow, D. Jung, and J. Jeon, "A big data based deep learning approach for vehicle speed prediction," in *2017 IEEE 26th International Symposium on Industrial Electronics (ISIE)*. IEEE, 2017, pp. 389–394.
- [6] H. Wang, Y. Cai, X. Chen, and L. Chen, "Night-time vehicle sensing in far infrared image with deep learning," *Journal of Sensors*, vol. 2016, 2016.
- [7] C. Kwan, B. Chou, A. Echavarren, B. Budavari, J. Li, and T. Tran, "Compressive vehicle tracking using deep learning," in *IEEE Ubiquitous Computing, Electronics & Mobile Communication Conference*, 2018.
- [8] D. K. Soother and J. Daudpoto, "A brief review of condition monitoring techniques for the induction motor," *Transactions of the Canadian Society for Mechanical Engineering*, vol. 43, no. 4, pp. 499–508, 2019.
- [9] Y. Jeong, S. Son, E. Jeong, and B. Lee, "An integrated self-diagnosis system for an autonomous vehicle based on an iot gateway and deep learning," *Applied Sciences*, vol. 8, no. 7, p. 1164, 2018.
- [10] —, "A design of a lightweight in-vehicle edge gateway for the self-diagnosis of an autonomous vehicle," *Applied Sciences*, vol. 8, no. 9, p. 1594, 2018.
- [11] J. P. N. González, L. E. G. Castañón, A. Rabhi, A. El Hajjaji, and R. Morales-Menendez, "Vehicle fault detection and diagnosis combining aann and anfi," *IFAC Proceedings Volumes*, vol. 42, no. 8, pp. 1079–1084, 2009.
- [12] J. P. N. González, "Fault diagnosis of a vehicle based on a history data hybrid approach," *Journal of Man, Machine and Technology, AICIT*, vol. 1, no. 1, pp. 63–72, 2012.
- [13] —, "Vehicle fault detection and diagnosis combining an aann and multiclass svm," *International Journal on Interactive Design and Manufacturing (IJIDeM)*, vol. 12, no. 1, pp. 273–279, 2018.
- [14] Y. Sun, W. Yu, Y. Chen, and A. Kadam, "Time series anomaly detection based on gan," in *2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS)*. IEEE, 2019, pp. 375–382.
- [15] J. Yin and W. Zhao, "Fault diagnosis network design for vehicle on-board equipments of high-speed railway: A deep learning approach," *Engineering Applications of Artificial Intelligence*, vol. 56, pp. 250–259, 2016.
- [16] S. S. Moosavi, A. Djerdir, Y. Ait-Amirat, D. A. Khaburi, and A. N'Diaye, "Artificial neural network-based fault diagnosis in the ac-dc converter of the power supply of series hybrid electric vehicle," *IET Electrical Systems in Transportation*, vol. 6, no. 2, pp. 96–106, 2016.
- [17] A. Glowacz, "Acoustic based fault diagnosis of three-phase induction motor," *Applied Acoustics*, vol. 137, pp. 82–89, 2018.
- [18] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung, "Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit," *Nature*, vol. 405, no. 6789, pp. 947–951, 2000.
- [19] Y. Gal, J. Hron, and A. Kendall, "Concrete dropout," in *Advances in neural information processing systems*, 2017, pp. 3581–3590.
- [20] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should i trust you?: Explaining the predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 2016, pp. 1135–1144.